Computing a Relevant Set of Nonbinary Maximum Acyclic Agreement Forests

Benjamin Albrecht *

Institut für Informatik, Ludwig-Maximilians-Universität, Germany

arXiv:1512.05703v1 [q-bio.PE] 17 Dec 2015

Abstract. There exist several methods dealing with the reconstruction of rooted phylogenetic networks explaining different evolutionary histories given by rooted binary phylogenetic trees. In practice, however, due to insufficient information of the underlying data, phylogenetic trees are in general not completely resolved and, thus, those methods can often not be applied to biological data. In this work, we make a first important step to approach this goal by presenting the first algorithm — called allMulMAAFs — that enables the computation of all relevant nonbinary maximum acyclic agreement forests for two rooted (nonbinary) phylogenetic trees on the same set of taxa. Notice that our algorithm is part of the freely available software Hybroscale computing minimum hybridization networks for a set of rooted (nonbinary) phylogenetic trees on an overlapping set of taxa.

Keywords Directed Acyclic Graphs · Hybridization · Agreement Forests · Bounded Search · Phylogenetics

* Benjamin Albrecht
Institut für Informatik, Ludwig-Maximilians-Universität, Germany
Tel.: +49-89-2180-4069
E-mail: albrecht@bio.ifi.lmu.de

1 Introduction

Evolution is often modeled by a rooted phylogenetic tree, which is a rooted tree whose edges are directed from its root to its leaf. The leaf set of such a tree is typically labeled by a set of species or strains of species and is in general called taxa set. Moreover, each internal node has in-degree one and out-degree two representing a certain speciation event. In applied phylogenetics, however, trees can contain nodes of out-degree larger or equal to three because, often, in order to resolve a certain ordering of speciation events, there is only insufficient information available and the common way to model this uncertainty is to use *nonbinary* nodes. This means, in particular, that *nonbinary* rooted phylogenetic trees contain several binary trees each each representing a different order of speciation events.

Note that there are two ways of interpreting a nonbinary node, which is also called *multifurcation* or *polytomy*. A multifurcation is *soft*, if it represents a lack of resolution of the true relationship as already mentioned above. A *hard* multifurcation, in contrast, represent a simultaneous speciation event of at least three species. However, as such events are assumed to be rare in nature, in this paper we consider all multifurcations of a pyhlogenetic tree as being soft.

A phylogenetic tree is usually built on homologous genes corresponding to a certain set of species. Thus, when constructing two phylogenetic trees based on two different homologous genes each belonging to the same set of species, one usually expects that both inferred trees are identical. Due to biological reasons, apart from other more technical reasons, however, these trees can be incongruent. One of those biological reasons are reticulation events that are processes including hybridization, horizontal gene transfer, or recombination. Now, based on such trees, one is interested to investigate the underlying reticulate evolution which can be done, for example, by computing rooted phylogenetic networks being structures similar to phylogenetic trees but, in contrast, can contain nodes of in-degree larger or equal to two. In a biological context, each of those nodes in such a network represents a certain reticulation event and, thus, a node with multiple incoming edges is called reticulation node or, in respect of hybridization, hybridization node.

Given two rooted binary phylogenetic trees T_1 and T_2 on the same set of taxa $\mathcal{X}$, a phylogenetic network, also called hybridization network, can be calculated by applying two major steps. In a first step, one can compute an agreement forest consisting of edge disjoint subtrees each being part of both input trees. Moreover, this set of subtrees $\mathcal{F}$ has to be a partition on $\mathcal{X}$ that is not allowed to have any conflicting ancestral relations in terms of T_1 and T_2. This means, in particular, for each pair of subtrees F_1 and F_2 in $\mathcal{F}$, if regarding T_1 the subtree corresponding to F_1 lies above or below the subtree corresponding to F_2, the same scenario has to be displayed in T_2. Notice that, by saying a subtree F_1 lies above F_2 we mean that there is a directed path leading from the root node corresponding to F_1 to the root node corresponding to F_2. Moreover, if this property holds for each pair of subtrees, we say that $\mathcal{F}$ is acyclic. Now, based on such an acyclic agreement forest, in a second step, one can apply a network construction algorithm, e.g., the algorithm HybridPhylogeny [6], which glues together each component by introducing reticulation nodes such that the resulting network displays both input trees. If such a network additionally provides a minimum number of reticulation events, it is called a minimum hybridization network. Notice that, as this is a highly combinatorial problem, for two rooted binary phylogenetic trees there typically exist multiple acyclic agreement forests and a hybridization network corresponding to one of those agreement forests is rarely unique.

Computing maximum acyclic agreement forests, which are acyclic agreement forests of minimum cardinality, is a NP-hard as well as APX-hard problem [8]. It has been shown, for the binary case, that, based on such a maximum acyclic agreement forest $\mathcal{F}$, there exists a minimum hybridization network whose number of provided reticulation events is one less than the number of components in $\mathcal{F}$ [5]. While, given two rooted binary phylogenetic trees, there exist some algorithms solving the maximum acyclic agreement forests as well as the minimum hybridization network problem, the nonbinary variant of both problems has attracted only less attention. More precisely, although there exist some approximation algorithms as well as exact algorithms [11, 16] computing maximum nonbinary agreement forests, until now there does not exist an algorithm computing nonbinary minimum hybridization networks.

The algorithm presented in this work was developed in respect to the algorithm allHNetworks [2, 3] computing a relevant set of minimum hybridization networks for multiple binary phylogenetic trees on the

same set of taxa. Broadly speaking, those networks are computed by inserting each of the input trees step wise to a so far computed network which is done, basically, in three major steps. In a first step, all embedded trees of a so far computed network N are extracted by selecting exactly one in-edge of each reticulation node. Notice that, at the beginning, N is initialized with the first input tree T_1 containing only one embedded tree. The more reticulation edges exist in a network, the more embedded trees can be extracted. More precisely, given r reticulation nodes of in-degree 2, there may exist up to 2^r different embedded trees. Now, for each of those embedded trees T', in a second step, all maximum acyclic agreement forests corresponding to T' and the current input tree T_i are computed which is done by applying the algorithm allMAAFs [14]. Finally, each component (except the root component F_ρ) of such a maximum acyclic agreement forest is inserted into the so far computed network in a certain way.

In this article, we present the first non-naive algorithm — called allMulMAAFs — calculating a relevant set of nonbinary maximum acyclic agreement forests for two rooted (nonbinary) phylogenetic trees on the same set of taxa, which is a first step to adapt the algorithm allHNetworks to trees containing nodes providing more than two outgoing edges. Our paper is organized as follows. We first introduce some basic definitions. Next, the algorithm allMulMAFs is presented that is extended in a subsequent section to the algorithm allMulMAAFs by introducing a tool through which agreement forests can be turned into acyclic agreement forests.

2 Preliminaries

In this section, we introduce some basic definitions referring to phylogenetic trees, hybridization networks, and nonbinary agreement forests based on the work of Huson *et al.* [9], Scornavacca *et al.* [14], and van Iersel *et al.* [11]. We therefor assume that the reader is familiar with basic graph-theoretic concepts.

Phylogenetic trees. A rooted phylogenetic $\mathcal{X}$-tree T is a tree whose edges are directed from the root to the leaves and whose nodes, except for the root, have a degree not equal to 2. There exists exactly one node of in-degree 0, namely the root. Each inner node has in-degree 1 and out-degree larger or equal to 2 whereas each leaf has in-degree 1 and out-degree 0. If T is a *bifurcating* or *binary* tree, its root has in-degree 0 and out-degree 2, each inner node an in-degree of 1 and an out-degree of 2, and each leaf an in-degree of 1 and out-degree 0. Notice that, in order to emphasize that a tree T can contain inner nodes of in-degree larger than 2, we say that T is a *multifurcating* or *nonbinary* tree. The leaves of a rooted phylogenetic $\mathcal{X}$-tree are bijectivley labeled by a taxa set $\mathcal{X}$, which usually consists of certain species or genes and is denoted by $\mathcal{L}(T)$. Considering a node v of T, the label set $\mathcal{L}(v)$ refers to each taxon that is contained in the subtree rooted at v. Additionally, given a set of phylogenetic trees $\mathcal{F}$, the label set $\mathcal{L}(\mathcal{F})$ denotes the union of each label set $\mathcal{L}(F_i)$ of each tree F_i in $\mathcal{F}$.

Now, based on a taxa set $\mathcal{X}' \subseteq \mathcal{X}$, we can define a restricted subtree of a rooted phylogenetic $\mathcal{X}$-tree, denoted by $T|_{\mathcal{X}'}$. The restricted tree $T|_{\mathcal{X}'}$ is computed by first repeatedly deleting each leaf that is either unlabeled or whose taxon is not contained in $\mathcal{X}'$, resulting in a subgraph denoted by $T(\mathcal{X}')$, and then by suppressing each node of both in- and out-degree 1.

A rooted nonbinary phylogenetic $\mathcal{X}$-tree T can contain *mulitfurcating* or *nonbinary* nodes, which are nodes of out-degree larger than or equal to 3. We say a rooted phylogenetic $\mathcal{X}$-tree T' is a *refinement* of T, if we can obtain T from T' by *contracting* some of its edges. More precisely, an edge $e = (u, v)$, with C_v being the set of children of v, is contracted by first deleting v together with all of its adjacent edges (including e) and then by reattaching each node c_i in C_v back to u by a inserting a new edge (u, c_i). Moreover, in this context we further say that T' is a *binary refinement of* T, if T' is binary. Similarly, if T' is a refinement of T, we can obtain T' from T by *resolving* some of its multifurcating nodes in the following way (cf. Fig. 1). Let v be a multifurcating node and let $C_v = \{c_1, \ldots, c_n\}$ be its set of children, then, we can resolve v as follows. First, a new node w is created, which is attached to v by inserting a new edge (v, w). Second, we select a subset C'_v of C_v, with $1 < |C'_v| < |C_v|$, and, finally, we prune each node c_i of C'_v from v and reattach c_i to w by inserting a new edge (w, c_i).

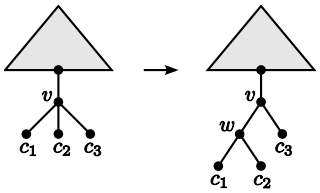

Fig. 1. Resolving a multifurcating node v by re-attaching c_1 and c_2 to a new inserted node w.

Now, let $\mathcal{X}'$ be a subset of the taxa set corresponding to a rooted phylogenetic $\mathcal{X}$-tree T. Then, the *lowest common ancestor* of T corresponding to $\mathcal{X}'$, shortly denoted by $\mathrm{LCA}_T(\mathcal{X}')$, is the farthest node v from the root in T with $\mathcal{X}' \subseteq \mathcal{L}(v)$. More precisely, v is chosen such that $\mathcal{X}' \subseteq \mathcal{L}(v)$ holds and there does not exist a node w with $\mathcal{X}' \subseteq \mathcal{L}(w)$ and $\mathcal{L}(w) \subset \mathcal{L}(v)$.

Phylogenetic networks. A *rooted phylogenetic network* N on $\mathcal{X}$ is a rooted connected digraph whose edges are directed from the root to the leaves as defined in the following. There is exactly one node of in-degree 0, namely the *root*, and no nodes of both in- and out-degree 1. The set of nodes of out-degree 0 is called the *leaf set of* N and is labeled one-to-one by the *taxa set* $\mathcal{X}$, also denoted by $\mathcal{L}(N)$. In contrast to a phylogenetic tree, such a network may contain undirected but not any directed cycles. Consequently, N can contain nodes of in-degree larger than or equal to 2, which are called *reticulation nodes* or *hybridization nodes*. Moreover, each edge that is directed into such a reticulation node is called *reticulation edge hybridization edge*.

Hybridization networks. A hybridization network N for a set of rooted nonbinary phylogenetic $\mathcal{X}$-trees $\mathcal{T}$ is a rooted phylogenetic network on $\mathcal{X}$ *displaying* a refinement T_i' of each tree T_i in $\mathcal{T}$. More precisely, this means that for each tree T_i in $\mathcal{T}$ there exists a set $E_i' \subseteq E(N)$ of reticulation edges *referring* to its refinement T_i'. This means, in particular, that we can obtain the tree T_i' from N by first deleting all reticulation edges that are not contained in E_i' and then suppress all nodes of both in- and out-degree 1. In this context, a reticulation edge (or hybridization edge) is an edge that is directed into a node with in-degree larger than or equal to 2, which is denoted as reticulation node (or hybridization node).

Given a hybridization network for a set of rooted nonbinary phylogenetic $\mathcal{X}$-trees $\mathcal{T}$, the *reticulation number* $r(N)$ is defined by

$$r(N) = \sum_{v \in V : \delta^-(v) > 0} (\delta^-(v) - 1) = |E| - |V| + 1, \tag{1}$$

where V refers to the set of nodes of N and $\delta^-(v)$ denotes the in-degree of a node v in V. Moreover, based on the definition of the reticulation number, the (minimum or exact) hybridization number $h(\mathcal{T})$ for $\mathcal{T}$ is defined by

$$h(\mathcal{T}) = \min\{r(N) : N \text{ displays a refinement of each } T_i \in \mathcal{T}\}. \tag{2}$$

A hybridization network displaying a set of rooted nonbinary phylogenetic $\mathcal{X}$-trees $\mathcal{T}$ with minimum hybridization number $h(\mathcal{T})$ is called a *minimum hybridization network*. Notice that even in the simplest case, if $\mathcal{T}$ consists only of two rooted binary phylogenetic $\mathcal{X}$-trees, the problem of computing the hybridization number is known to be *NP-hard* but fixed-parameter tractable [7, 8], which means that the problem is exponential in some parameter related to the problem itself, namely the hybridization number of $\mathcal{T}$, but only at most polynomial in its input size, which is, in this context, the number of nodes and edges in $\mathcal{T}$.

Forests. Let T be a rooted nonbinary phylogenetic $\mathcal{X}$-tree T. Then, we call any set of rooted nonbinary phylogenetic trees $\mathcal{F} = \{F_1, \ldots, F_k\}$ with $\mathcal{L}(\mathcal{F}) = \mathcal{X}$ a *forest on* $\mathcal{X}$, if we have for each pair of trees F_i and F_j that $\mathcal{L}(F_i) \cap \mathcal{L}(F_j) = \emptyset$. Moreover, we say that $\mathcal{F}$ is a *forest for* T, if additionally for each component F in $\mathcal{F}$ the tree F is a refinement of $T|_{\mathcal{L}(F)}$. Lastly, given two forests $\mathcal{F}$ and $\hat{\mathcal{F}}$ for a rooted phylogenetic $\mathcal{X}$-tree T, we say that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$, if for each component $\hat{F}$ in $\hat{\mathcal{F}}$ there exists a component F in

$\mathcal{F}$ such that $\hat{F}$ is a binary resolution of F. Lastly, let $\mathcal{F}$ be a forest on a taxa set $\mathcal{X}$, then by $\overline{\mathcal{F}}$ we refer to the forest that is obtained from $\mathcal{F}$ by deleting each element only consisting of an isolated node.

Nonbinary agreement forests. Given two rooted nonbinary phylogenetic $\mathcal{X}$-trees T_1 and T_2. For technical purpose, we consider the root of both trees T_1 and T_2 as being a node that has been marked by new taxon $\rho \notin \mathcal{X}$. More precisely, let r_i be the root of the tree T_i with $i \in \{1, 2\}$. Then, we first create a new node v_i as well as a new leaf ℓ_i labeled by a new taxon $\rho \notin \mathcal{X}$ and then attach these nodes to r_i by inserting the two edges (v_i, r_i) and (v_i, ℓ_i) such that v_i is the new root of the resulting tree. Now, an *agreement forest* for two so marked trees T_1 and T_2 is a forest $\mathcal{F} = \{F_\rho, F_1, \ldots, F_k\}$ on $\mathcal{X} \cup \{\rho\}$ satisfying the following three conditions.

(1) Each component F_i with taxa set $\mathcal{X}_i$ equals a refinement of $T_1|_{\mathcal{X}_i}$ and $T_2|_{\mathcal{X}_i}$, respectively.
(2) There is exactly one component, denoted as F_ρ, containing ρ.
(3) Let $\mathcal{X}_\rho, \mathcal{X}_1, \ldots, \mathcal{X}_k$ be the taxa sets corresponding to $F_\rho, F_1, \ldots, F_k$. All trees in $\{T_1(\mathcal{X}_i) | i \in \{\rho, 1, \ldots, k\}\}$ and $\{T_2(\mathcal{X}_i) | i \in \{\rho, 1, \ldots, k\}\}$ are edge disjoint subtrees of T_1 and T_2, respectively.

Let $\mathcal{F}$ be an agreement forests for two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2 and let E be a set only consisting of edges in $\mathcal{F}$. Then, by $\mathcal{F} \ominus E$ we refer to the forest $\mathcal{F}'$ that is obtained from $\mathcal{F}$ by contracting each edge in E. Based on this definition, we say that $\mathcal{F}$ is *relevant*, if there does not exist an edge e in $\mathcal{F}$ such that $\mathcal{F} \ominus \{e\}$ is still an agreement forest for T_1 and T_2.

Moreover, a *maximum agreement forest* for two rooted nonbinary phylogenetic $\mathcal{X}$-trees T_1 and T_2 is an agreement forest of minimal size, which implies that there does not exist a smaller set of components fulfilling the properties of an agreement forest for T_1 and T_2 listed above. Additionally, we call an agreement forest $\mathcal{F}$ for T_1 and T_2 *acyclic*, if its underlying *ancestor-descendant graph* $AG(T_1, T_2, \mathcal{F})$ does not contain any directed cycles (cf. Fig. 2). This directed graph contains one node corresponding to precisely one component of $\mathcal{F}$ and an edge (F_i, F_j) for a pair of its nodes F_i and F_j, with $i \neq j$, if,

(i) regarding T_1, there is a path leading from the root of $T_1(\mathcal{X}_i)$ to the root of $T_1(\mathcal{X}_j)$ containing at least one edge of $T_1(\mathcal{X}_i)$,
(ii) or, regarding T_2, there is a path leading from the root of $T_2(\mathcal{X}_i)$ to the root of $T_2(\mathcal{X}_j)$ containing at least one edge of $T_2(\mathcal{X}_i)$.

In this context, $\mathcal{X}_i \subseteq \mathcal{X}$ and $\mathcal{X}_j \subseteq \mathcal{X}$ refers to the set of taxa that are contained in F_i and F_j, respectively. Again, we call an acyclic agreement forest consisting of a minimum number of components a *maximum acyclic agreement forest*. Notice that, in the binary case, for a maximum acyclic agreement forest containing k components there exists a hybridization network whose reticulation number is $k - 1$ [5]. This means, in particular, if a maximum acyclic agreement forest for two binary phylogenetic $\mathcal{X}$-trees contains only one component, both trees are congruent.

Acyclic orderings. Now, if $\mathcal{F}$ is acyclic and, thus, $\mathrm{AG}(T_1, T_2, \mathcal{F})$ does not contain any directed cycles, one can compute an acyclic ordering Π as already described in the work of Baroni *et al.* [6]. First, select the node v_ρ corresponding to F_ρ of in-degree 0 and remove v_ρ together with all its incident edges. Next, again choose a node v_1 with in-degree 0 and remove v_1. By continuing this way, until all nodes have been removed, one receives the ordering $\Pi = (v_\rho, v_1, \ldots, v_k)$. Notice that, since the graph does not contain any cycles, such an ordering always has to exist. In the following, we call the ordering of components corresponding to each node in Π, denoted by $(F_\rho, F_1, \ldots, F_k)$, an acyclic ordering of $\mathcal{F}$. As during each of those steps there can occur multiple nodes of in-degree 0, especially, if $\mathcal{F}$ contains multiple components consisting only of one taxon, such an acyclic ordering is in general not unique.

Trees reflecting agreement forests. Let $\mathcal{F} = \{F_\rho, F_1, F_2, \ldots, F_{k-1}\}$ be an agreement forest for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, then, a tree $T_i(\mathcal{F})$ for $i \in \{1, 2\}$ corresponds to the tree T_i reflecting each component in $\mathcal{F}$ (cf. Fig. 3). Generally speaking, in such a tree some of its nodes are

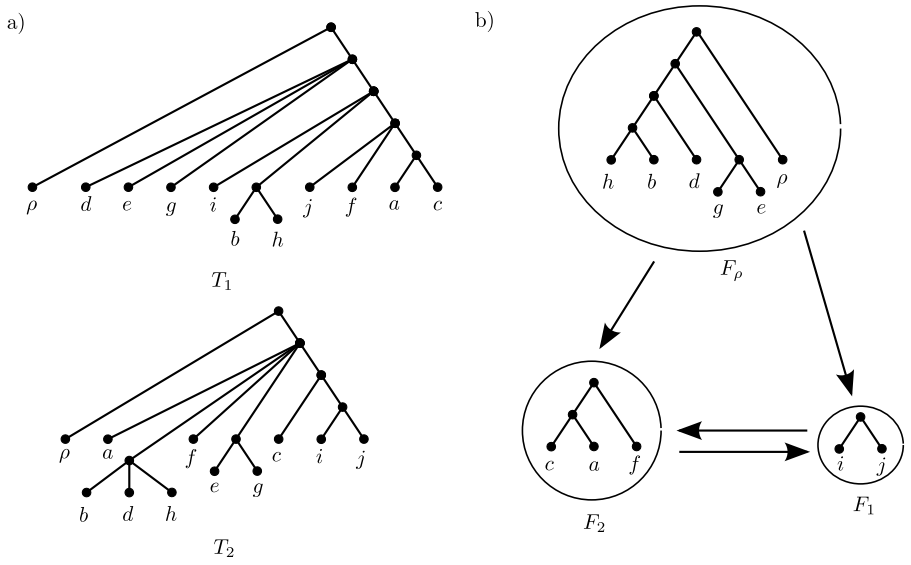

Fig. 2. (a) Two rooted nonbinary phylogenetic $\mathcal{X}$-trees T_1 and T_2. (b) The ancestor-descendant graph $AG(T_1, T_2, \mathcal{F})$ with $\mathcal{F} = \{F_\rho, F_1, F_2\}$. Notice that the component corresponding to each node of the graph is drawn inside.

resolved such that the definition of an agreement forest can be applied in terms of the resulting tree and $\mathcal{F}$. Technically speaking, such a tree $\hat{T}_i = T_i(\mathcal{F})$ satisfies the following two properties.

(1) Each component F_j in $\mathcal{F}$ refers to a restricted subtree of $\hat{T}_i|_{\mathcal{L}(F_j)}$.
(2) All trees in $\{\hat{T}_i(\mathcal{L}(F_j))|j \in \{\rho, 1, \ldots, k-1\}\}$ are node disjoint subtrees in $\hat{T}_i$.

We can construct such a tree $T_i(\mathcal{F})$ by reattaching the components of $\mathcal{F}$ back together in a specific way as follows. Let $\Pi_\mathcal{F} = (F_0, F_1, \ldots, F_k)$, with $F_0 = F_\rho$, be an acyclic ordering that can be obtained from $AG(T_i, T_i, \mathcal{F})$ as discussed above. Notice that, as this graph is based only on one of both trees, this graph cannot contain any directed cycles and, thus, $\Pi_\mathcal{F}$ always exists. Now, each of those components in $\Pi_\mathcal{F}$, beginning with F_1, is added sequentially to a growing tree T^* (initialized with F_0) as follows.

(i) Let $\mathcal{X}_{<m}$ be the union of each taxa set corresponding to each component F_l in $\Pi_\mathcal{F}$ with $l < m$, i.e., $\mathcal{X}_{<m} = \bigcup_l^{m-1} \mathcal{L}(F_l)$, and let $\mathcal{X}_m$ be the taxa set corresponding to F_m. Moreover, let $P_m = (v_0^m, v_1^m, \ldots, v_n^m)$ be those nodes lying on the path connecting the node $v_0^m = \mathrm{LCA}_{T_i}(\mathcal{X}_m)$ and the root v_n^m of T_i such that v_q^m, with $q \in \{1, \ldots, n\}$, is the parent of v_{q-1}^m. Then,

$$v' = \min_q \{v_q^m : v_q^m \in P \wedge \mathcal{L}(v_q^m) \cap \mathcal{X}_{<m} \neq \emptyset\}.$$

(ii) Let $\mathcal{X}'$ be the set of taxa corresponding to the leaf set of $T_i(v')$ restricted to $\mathcal{X}_{<m}$. Notice that, due to the definition of v', this set $\mathcal{X}'$ is not empty. Moreover, based on $\mathcal{X}'$, let v^* be the node in T^* corresponding to $\mathrm{LCA}_{T^*}(\mathcal{X}')$.
(iii) Now, given v^*, the component F_m is added to T^* by connecting its root node ρ_m to the in-edge of v^*. More precisely, first a new node x is inserted into the in-edge of v^* and then ρ_m is connected to x by inserting a new edge (x, ρ_m).

Notice that, since there can exist multiple acyclic orderings for an acyclic agreement forest $\mathcal{F}$, the tree $T_i(\mathcal{F})$ is in general not unique.

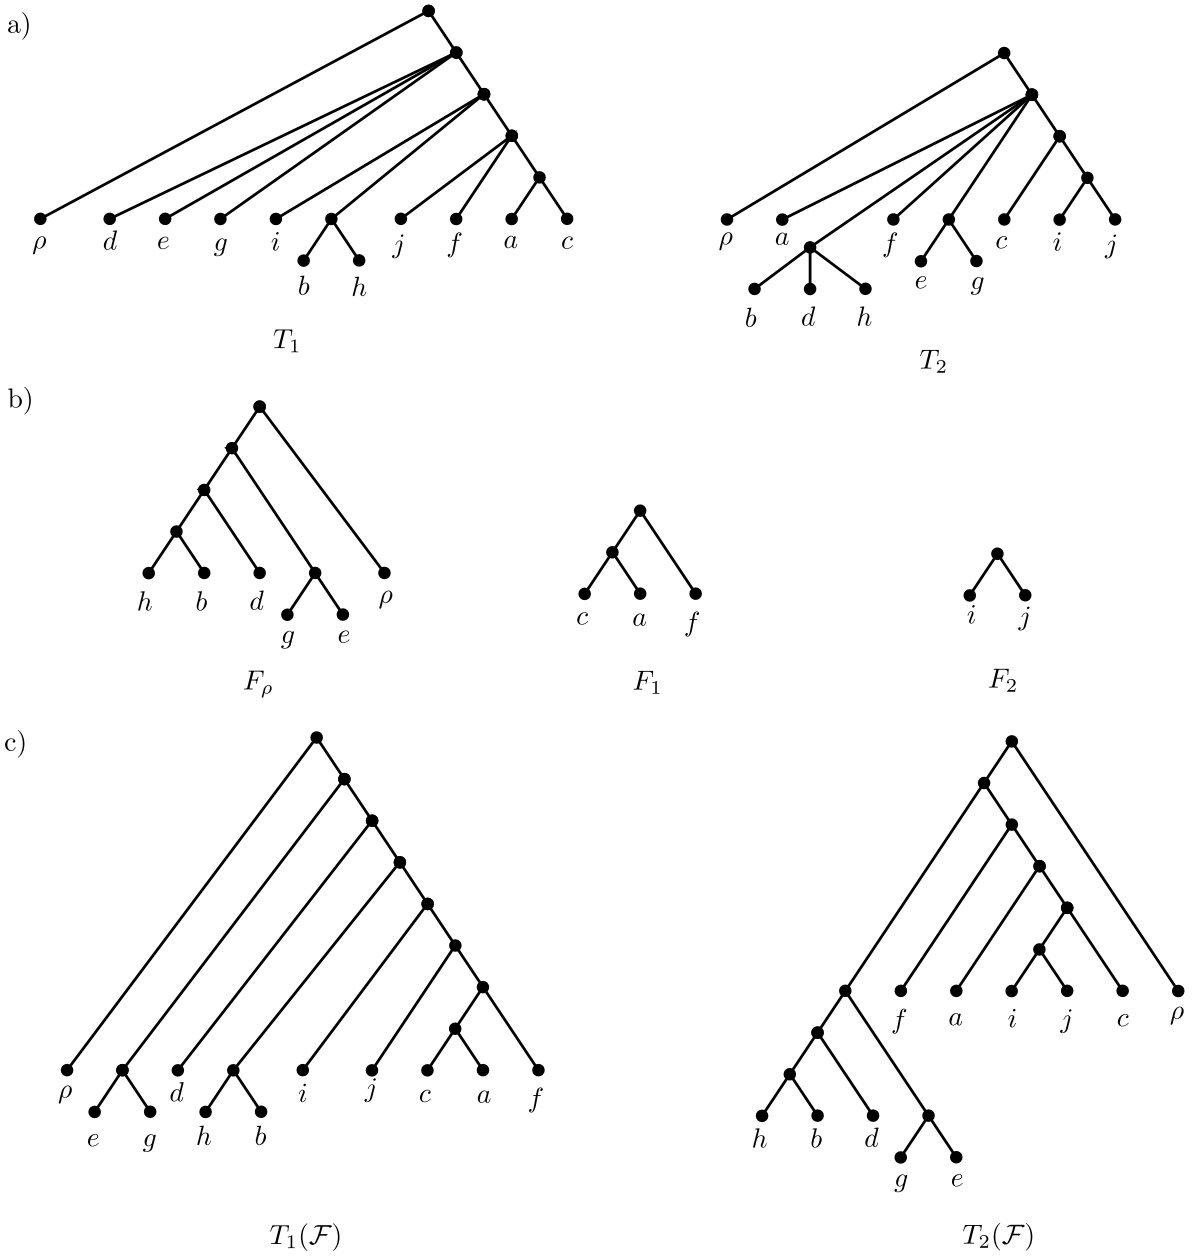

Fig. 3. (a) Two rooted nonbinary phylogenetic $\mathcal{X}$-trees T_1 and T_2. (b) An agreement forest $\mathcal{F}$ for T_1 and T_2. (c) Two trees $T_1(\mathcal{F})$ and $T_2(\mathcal{F})$ both reflecting $\mathcal{F}$ and being calculated in terms of the acyclic ordering (F_ρ, F_2, F_1) and (F_ρ, F_1, F_2), respectively.

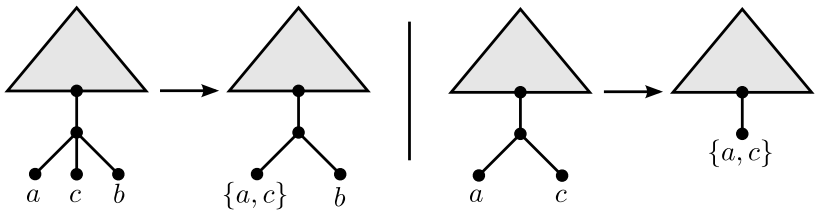

Fig. 4. Two different ways of contracting a cherry $\{a, c\}$ depending on its number of siblings which is one on the left hand side and zero on the right hand side and .

3 The algorithm allMulMAFs

In this section, we show how to modify the algorithm allMAAFs [14] such that the output consists of all relevant maximum agreement forests for two rooted *nonbinary* phylogenetic $\mathcal{X}$-trees. Similar to the algorithm allMAAFs, this algorithm is again based on processing common and contradicting cherries. In order to cope with nonbinary nodes, however, now for an internal node one has to consider more than one cherry and, before cutting a particular set of edges, one first has to resolve some nonbinary nodes. Furthermore, in respect to the definition of relevant maximum agreement forests, when expanding contracted nodes one has to take care on not generating any contractible edges.

In the following, we will first introduce some further notations necessary for describing the algorithm allMulMAFs. Moreover, we give a detailed formal proof establishing the correctness of the algorithm, which means, in particular, that we will show that the algorithm calculates all relevant maximum agreement forests for two rooted *nonbinary* phylogenetic $\mathcal{X}$-trees. Finally, we end this section by discussing its theoretical worst-case runtime.

3.1 Notations

Before going into details, we have to give some further notations that are crucial for the following description of the algorithm.

Removing leaves. Given a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree R, a *leaf ℓ is removed* by first deleting its in-edge and then by suppressing its parent p, if, after ℓ has been deleted, p has out-degree 1.

Cherries. Let R be a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree and let ℓ_a and ℓ_c be two of its leaves that are adjacent to the same parent node p and labeled by taxon a and c, respectively. Then, we call the set consisting of the two taxa $\{a, c\}$ a *cherry of R*, if the children of p are all leaves. Now, let $\{a, c\}$ be a cherry of R and let $\mathcal{F}$ be a forest on a taxa set $\mathcal{X}'$ such that $\overline{\mathcal{F}}$ is a forest for R. Then, we say $\{a, c\}$ is a *contradicting cherry of R and $\mathcal{F}$*, if $\mathcal{F}$ does not contain a tree containing $\{a, c\}$. Otherwise, if such a tree exists in $\mathcal{F}$, the cherry $\{a, c\}$ is called a *common cherry of R and $\mathcal{F}$*.

Contracting cherries. Given a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree R, a cherry $\{a, c\}$ of R can be contracted in two different ways (cf. Fig. 4). Either, if the two leaves ℓ_a and ℓ_c labeled by a and c, respectively, are the only children of its parent node p, first the in-edge of both nodes is deleted and then the label of p is set to $\{a, c\}$. Otherwise, if the two leaves ℓ_a and ℓ_c contain further siblings and, thus, its parent p has an out-degree larger than 2, just the in-edge of ℓ_a is deleted and the label of ℓ_c is replaced by $\{a, c\}$.

Cutting edges. Given a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree F, the in-edge e_v of a node v is cut as follows (cf. Fig. 5). First e_v is deleted and then its parent node p is suppressed, if, after the deletion of e_v, p has out-degree 1. Note that by cutting an edge in F, two rooted (nonbinary) phylogenetic trees with taxa set $\mathcal{X}'$ and $\mathcal{X}''$ are generated with $\mathcal{X}' \cup \mathcal{X}'' = \mathcal{X}$ and $\mathcal{X}' \cap \mathcal{X}'' = \emptyset$.

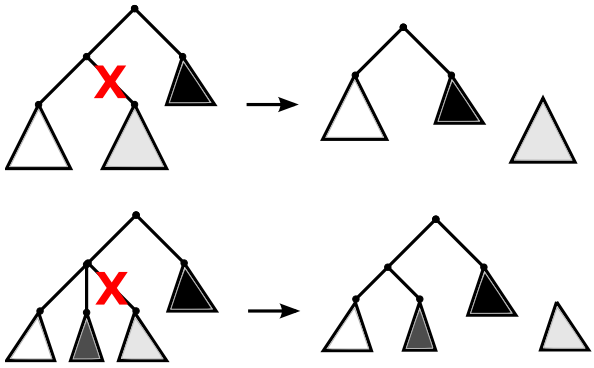

Fig. 5. Two different ways of cutting an edge depending on the out-degree of its source node which is 2 at the top and 3 at the bottom.

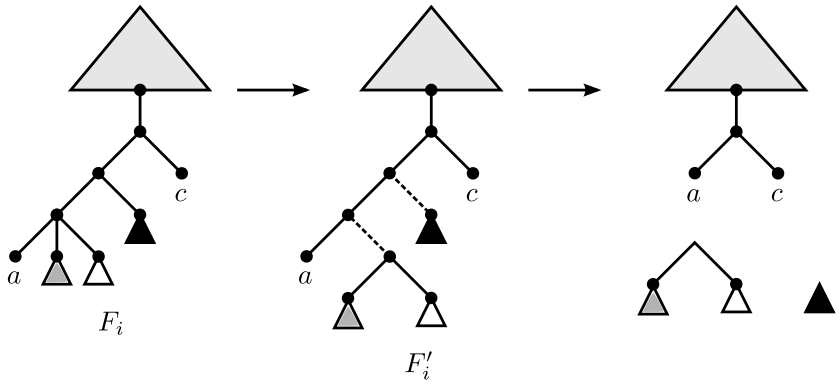

Fig. 6. An illustration of Case 3b. First F_i is refined into $F'_i = F_i[a \sim c]$ and then each pendant subtree of the path connecting both nodes labeled by a and c, respectively, are cut. Note that each dashed edge of F'_i is part of the pendant edge set E_B for a and c.

Moreover, let $\mathcal{F}$ be a set of rooted (nonbinary) phylogenetic $\mathcal{X}$-trees and let E be a set of edges in which each edge e is part of a tree in $\mathcal{F}$. Then, in order to ease reading, we write $\mathcal{F} - E$ to denote the cutting of each edge $e \in E$ within its corresponding tree in $\mathcal{F}$.

Pendant edges. Given a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree F and two leaves ℓ_a and ℓ_c labeled by taxon a and c, respectively, that are not adjacent to the same node. Then, the set of *pendant edges* E_B *for* a *and* c is based on a refinement of F, shortly denoted by $F[a \sim c]$, which is obtained from F as follows. Let $(\ell_a, v_1, v_2, \ldots, v_n, \ell_c)$ be the path connecting the two leafs ℓ_a and ℓ_c in F. Then, each node $v \in \{v_1, v_2, \ldots, v_n\}$ with $\delta^+(v) > 2$ is turned into a node of out-degree 2 as follows. First a new node w is created that is attached to v by inserting a new edge (v, w) and then each out-going edge (v, x) with $x \notin \{w, \ell_a, v_1, v_2, \ldots, v_n, \ell_c\}$ is deleted followed by reattaching the node x to w by inserting a new edge (w, x). Now, regarding $F[a \sim c]$, let $(\ell'_a, v'_1, v'_2, \ldots, v'_n, \ell'_c)$ be the path connecting the two nodes ℓ'_a and ℓ'_c labeled by a and c, respectively. Then, E_B consists of each edge (u', v') with $u' \in \{v'_1, v'_2, \ldots, v'_n\}$ and $v' \notin \{\ell'_a, v'_1, v'_2, \ldots, v'_n, \ell'_c\}$ (cf. Fig. 6).

Moreover, given a forest $\mathcal{F}$ on $\mathcal{X}$ containing a tree F with two leaves ℓ_a and ℓ_c labeled by taxon a and c, respectively, that are not adjacent to the same node, then, by $\mathcal{F}[a \sim c]$ we refer to $\mathcal{F}$ in which F is replaced by $F[a \sim c]$.

Labeled nodes. Let R be a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree, then, by $\ell(R)$ we denote the number of its labeled nodes. Moreover, let $\mathcal{F}$ be a forest on $\mathcal{X}'$ such that $\overline{\mathcal{F}}$ is a forest for R. Then, $\ell(\overline{\mathcal{F}})$ refers to the

number of labeled nodes that are contained in each tree of $\overline{\mathcal{F}}$. Additionally, we write $\ell(R) \equiv \ell(\overline{\mathcal{F}})$, if $\ell(R)$ equals $\ell(\overline{\mathcal{F}})$ and if for each labeled node v_R in R there exists a labeled node v_F in $\overline{\mathcal{F}}$ such that $\mathcal{L}(v_R) = \mathcal{L}(v_F)$. Moreover, if $\ell(R) \equiv \ell(\overline{\mathcal{F}})$ holds, we say that a leaf ℓ in R *refers* to a leaf ℓ' in $\mathcal{F}$, if $\mathcal{L}(\ell)$ equals $\mathcal{L}(\ell')$.

3.2 The algorithm

In this section, we give a description of the algorithm allMulMAFs calculating a particular set of nonbinary maximum agreement forests for two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2. More specifically, as shown by an upcoming formal proof, this set consists of all relevant agreement forests for both trees. Before that, however, we want to give a remark emphasizing that the algorithm is based on a previous published algorithm that solves a similar problem dealing with rooted binary phylogenetic $\mathcal{X}$-trees.

Remark 1. Our algorithm is an extension of the algorithm allMAAFs [14] computing all maximum *acyclic* agreement forests for two rooted binary phylogenetic $\mathcal{X}$-trees T_1 and T_2. Notice that the work of Scornavacca *et al.* [14] also contains a formal proof showing the correctness of the presented algorithm. The algorithm allMulMAFs presented here has a similar flavor and, thus, our notation basically follows the notation that has already been used for the description of the algorithm allMAAFs.

Broadly speaking, given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2 as well as a parameter $k \in \mathbb{N}$, our algorithm acts as follows. Based on the topology of the first tree T_1, the second tree T_2 is cut into several components until either the number of those components exceeds k or the set of components fulfills each property of an agreement forest for T_1 and T_2. To ensure that there does not exist an agreement forest consisting of less than k components, the following steps can be simply conducted by step-wise increasing parameter k beginning with $k = 0$. Thus, as far as our algorithm reports an agreement forest for T_1 and T_2 of size k, this agreement forest must be of minimum size and, hence, must be a maximum agreement forest for both input trees. In order to speed up computation, one can either set k to a lower bound calculated by particular approximation algorithms as, for instance, given in van Iersel *et al.* [11], or directly to the hybridization number calculated by applying less complex algorithms, e.g., the algorithm TerminusEst [13].

The algorithm allMulMAFs takes as input two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2 as well as a parameter $k \in \mathbb{N}$. If $k < h(T_1, T_2)$ holds, an empty set is returned. Otherwise, as we will show later in Section 3.3, if k is larger than or equal to $h(T_1, T_2)$, the output $\boldsymbol{\mathcal{F}}$ of allMulMAFs contains all relevant maximum agreement forests for T_1 and T_2. Throughout the algorithm three specific tree operations are performed on both input trees. Either a leaf is removed, subtrees are cut, or a common cherry is contracted.

The algorithm allMulMAFs contains a recursive subroutine, in which the input of each recursion consists of a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree R, a forest $\mathcal{F}$ on some taxa set $\mathcal{X}'$ with $\overline{\mathcal{F}}$ being a forest for R, a parameter $k \in \mathbb{N}$, and a map M. This map M is necessary for undoing each cherry reduction that has been applied to each component of the resulting forest. For that purpose, M maps a set of taxa $\tilde{\mathcal{X}}$ to a triplet $(\mathcal{X}_1, \mathcal{X}_2, B)$ with $\mathcal{X}_1 \cup \mathcal{X}_2 = \tilde{\mathcal{X}}$, $\mathcal{X}_1 \cap \mathcal{X}_2 = \emptyset$, and $B \in \{\top, \bot\}$, where B denotes the way of how a cherry is expanded (as discussed below). In order to ease reading, by $M[\tilde{\mathcal{X}}] \leftarrow B$ we refer to the operation on M mapping $\tilde{\mathcal{X}}$ to B. This means, in particular, if M already contains an element with taxa set $\tilde{\mathcal{X}}$ this element is replaced. For instance, if $M[\tilde{\mathcal{X}}] = (\mathcal{X}_1, \mathcal{X}_2, \top)$, by $M[\tilde{\mathcal{X}}] \leftarrow \bot$ the taxa set $\tilde{\mathcal{X}}$ is remapped to $(\mathcal{X}_1, \mathcal{X}_2, \bot)$ so that after this operation $M[\tilde{\mathcal{X}}] = (\mathcal{X}_1, \mathcal{X}_2, \bot)$.

Expanding agreement forests. The expansion of an agreement forest $\mathcal{F}$ is done by applying the following steps to $\mathcal{F}$. Choose a leaf v corresponding to a component F_i in $\mathcal{F}$ whose taxon $\mathcal{L}(v)$ is contained in M. Let $(\mathcal{X}_1, \mathcal{X}_2, B)$ be the triplet referring to $M[\mathcal{L}(v)]$, then, depending on $B \in \{\top, \bot\}$, one of the following two operations is performed as illustrated in Figure 7.

- If B equals $\bot$, replace v in F_i by first creating two new nodes w_1 and w_2 and then by labeling w_1 and w_2 by $\mathcal{X}_1$ and $\mathcal{X}_2$, respectively. Finally, both nodes w_1 and w_2 are attached to v by inserting a new edge (v, w_1) and (v, w_2)

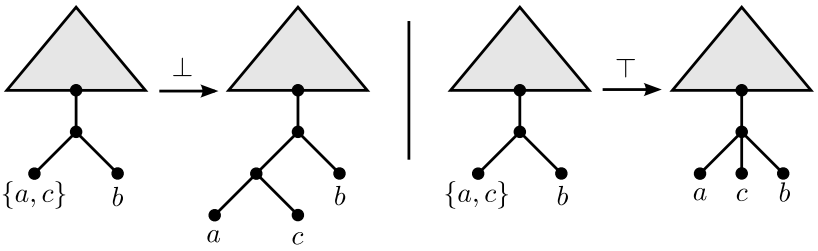

Fig. 7. Expanding a cherry $\{a, c\}$ in respect of $\perp$ (left) and in respect of $\top$ (right).

– Otherwise, if B equals $\top$, replace v in F_i by first creating a new node w and then by labeling w and v by $\mathcal{X}_1$ and $\mathcal{X}_2$, respectively. Finally, w is attached to the parent p of v by inserting a new edge (p, w).

These steps are repeated in an exhaustive way until each taxon of each leaf in $\mathcal{F}$ is not contained in M. As a result, the expanded forest corresponds to a nonbinary agreement forest for the two input trees T_1 and T_2. Notice that in the following, by saying a cherry is expanded *in respect of* $\perp$ or *in respect of* $\top$, we refer to one of both ways as described above.

In the following, a description of the recursive algorithm allMulMAFs is given. Here we assume that, at the beginning, R is initialized by T_1, $\mathcal{F}$ by $\{T_2\}$ and M by $\emptyset$, with T_1 and T_2 being two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees. Then, during each recursive call, $\overline{\mathcal{F}}$ is a forest for R with $\ell(R) \equiv \ell(\overline{\mathcal{F}})$ and, depending on the size of $\mathcal{F}$ (cf. Case 1a–c) and the choice of the next cherry that is selected from R (cf. Case 2 and 3), the following steps are performed.

Case 1a. If $\mathcal{F}$ contains more than k components, the computational path is aborted immediately and the empty set is returned.

Case 1b. If R only consists of a single leaf, each F_i in $\mathcal{F}$ is expanded as prescribed in M, and, finally, returned.

Case 1c. If there exists a specific leaf ℓ in R that refers to an isolated node in $\mathcal{F}$, this leaf ℓ is removed from R resulting in R'. Next, the algorithm branches into a new path by recursively calling the algorithm with R', $\mathcal{F}$, k, and M' corresponding to M where $\mathcal{L}(\ell)$ is re-mapped to $M[\mathcal{L}(\ell)] \leftarrow \perp$.
Otherwise, if such a leaf ℓ does not exist continue with Case 2.

Case 2. If there exists a common cherry $\{a, c\}$ of R and $\mathcal{F}$, the cherry $\{a, c\}$ is contracted in R and $\mathcal{F}$ resulting in R' and $\mathcal{F}'$. Second, the algorithm branches into a new path by recursively calling the algorithm with R', $\mathcal{F}'$, k, and M', where M' corresponds to M that has been updated as follows. If both parents of $\{a, c\}$ in R and $\mathcal{F}$ have out-degree ≥ 3, $\mathcal{L}(a) \cup \mathcal{L}(c)$ is mapped to $(\mathcal{L}(a), \mathcal{L}(c), \top)$, otherwise, $\mathcal{L}(a) \cup \mathcal{L}(c)$ is mapped to $(\mathcal{L}(a), \mathcal{L}(c), \perp)$.
Otherwise, if such a common cherry does not exist, continue with Case 3.

Case 3. If there does not exist a common cherry of R and $\mathcal{F}$, a node v in R whose children are all leaves is selected. Now, *for each* cherry $\{a, c\}$ of v, depending on the location of the leaves referring to a and c in $\mathcal{F}$, one of the following two cases is performed.

Case 3a. If $a \not\sim_{\mathcal{F}} c$ holds, and, thus, the leaves referring to a and c in $\mathcal{F}$ are located in two different components, the algorithm branches into two computational paths by recursively calling the algorithm by R, $\mathcal{F} - \{e_a\}$, k, and M' as well as R, $\mathcal{F} - \{e_c\}$, k, and M', where e_a and e_c correspond to the in-edge of the leaf of $\mathcal{F}$ referring to a and c, respectively, and M' is obtained from M as follows. Let p be the parent in $\mathcal{F}$ of the leaf referring to a (resp. c). If p has out-degree larger than 2, nothing is done. Otherwise, if p has out-degree 2, let

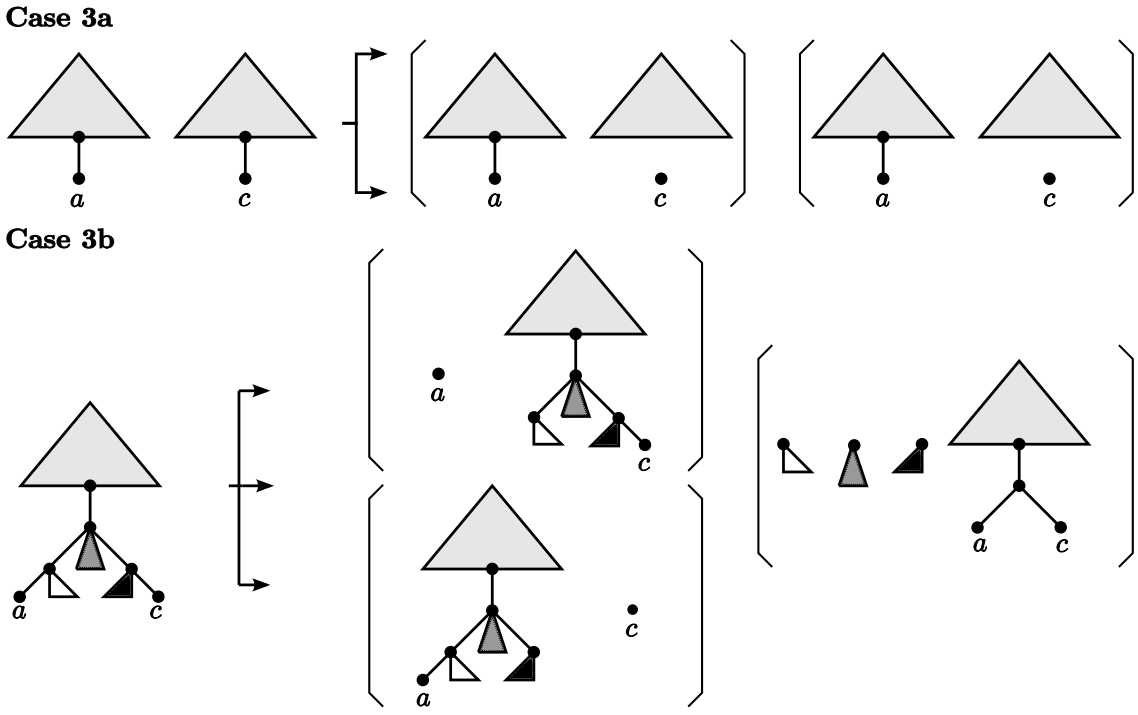

Fig. 8. An illustration of Case 3a branching into two computational paths and Case 3b branching into three computational paths. Regarding Case 3b, one additional computational path is created in which all pendant subtrees lying on the path connecting the two nodes labeled by a and c, respectively, are cut.

ℓ be the sibling of the leaf labeled by a (resp. c). Then, if ℓ is a leaf M is updated so that $M' = M[\mathcal{L}(\ell)] \leftarrow \bot$.

Case 3b. If $a \sim_{\mathcal{F}} c$ holds, and, thus, in $\mathcal{F}$ both leaves ℓ_a and ℓ_c referring to a and c, respectively, are located in the same component F_i, the algorithm branches into the following three computational paths. First, similar to Case 3a, the algorithm is called by R, $\mathcal{F} - \{e_a\}$, k, and M' as well as R, $\mathcal{F} - \{e_c\}$ k, and M'. Second, a third computational path is initiated by calling the algorithm with R, $\mathcal{F}[a \sim c] - E_B$, k, and M'', where E_B refers to the set of pendant edges in $\mathcal{F}[a \sim c]$ and M'' is obtained from M as follows.

Let $(\ell_a, v_1, \ldots, v_n, \ell_c)$ denote the path connecting ℓ_a and ℓ_c in $\mathcal{F}$. Then, M'' is obtained by updating M as follows. If v_1 does not correspond to $\mathrm{LCA}_{F_i}(\{a, c\})$, $\mathcal{L}(\ell_a)$ is remapped to $M[\mathcal{L}(\ell_a)] \leftarrow \bot$. Similarly, if v_n does not correspond to $\mathrm{LCA}_{F_i}(\{a, c\})$, $\mathcal{L}(\ell_c)$ is remapped $M[\mathcal{L}(\ell_c)] \leftarrow \bot$.

An illustration of this case is given in Figure 6

We end the description of the algorithm by noting that the algorithm allMulMAFs always terminates, since during each recursive call either the size of R decreases or the number of components in $\mathcal{F}$ increases. More precisely, the size of R is decreased by one either by deleting one of its leaves ℓ referring to an isolated node in $\mathcal{F}$ (cf. Case 1c) or by contracting a common cherry of R and $\mathcal{F}$ (cf. Case 2). If R is not decreased, at least one edge in $\mathcal{F}$ is cut (cf. Case 3) and, thus, its size increases at least by one. As each computational path of the algorithm stops if R only consists of an isolated node or if k edges have been cut, each recursive call does always make progress towards one of both abort criteria.

3.3 Correctness of allMulMAFs

In this section, we establish the correctness of the algorithm allMulMAFs. However, before doing so, we want to give an important remark emphasizing the relation of our algorithm allMulMAFs to the previously

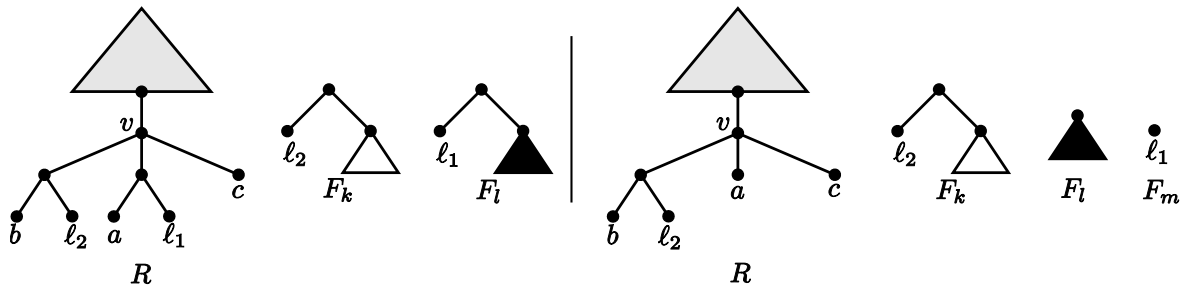

Fig. 9. (left) An illustration of a pseudo cherry $\{a, c\}$. (right) The result of preparing the pseudo cherry $\{a, c\}$ given on the left hand side.

presented algorithm allMAAFs^1 [4], which is a modification of the algorithm allMAAFs [14] improving the processing of contradicting cherries.

Remark 2. Given two binary phylogenetic $\mathcal{X}$-trees, the algorithm allMulMAFs processes an ordered set of cherries in the same way as the algorithm allMAAFs^1 omitting its acyclic check (henceforth denoted as allMAFs^1) testing an agreement forests for acyclicity. This means, in particular, that our algorithm allMulMAFs is simply an extension of the algorithm allMAFs^1 that is now able to handle nonbinary trees, but for binary trees still acts in the same way.

As a consequence of Remark 2, the upcoming proof showing the correctness of allMulMAFs refers to the correctness of allMAAFs^1 calculating all maximum acyclic agreement forests for two rooted binary phylogenetic $\mathcal{X}$-trees [4, Theorem 2]. In a first step, however, in order to ease the understanding of our proof, we will introduce a connective element between both algorithms, which is a modified version of our original algorithm — called allMulMAFs* — processing types of cherries that are not considered by computational paths corresponding to allMulMAFs.

3.3.1 The algorithm allMulMAFs* Before describing the algorithm, we have to add further definitions that are crucial for what follows.

Proper leaves. Given a leaf ℓ of a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree R labeled with taxon a as well as a forest $\mathcal{F}$ on some taxa set $\mathcal{X}'$ such that $\overline{\mathcal{F}}$ is a forest for R, ℓ is called a *proper leaf of* R and $\mathcal{F}$, if the corresponding leaf in $\mathcal{F}$ labeled by taxon a is a child of some root.

Pseudo cherries. Given a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree R as well as a forest $\mathcal{F}$ on some taxa set $\mathcal{X}'$ such that $\overline{\mathcal{F}}$ is a forest for R, we call a set of two taxa $\{a, c\}$ a *pseudo cherry for* R *and* $\mathcal{F}$, if the following two properties hold. First for each child v of $\mathrm{LCA}_R(\{a, c\})$ its leaf set $\mathcal{L}(T(v))$ of size n contains at least $n - 1$ proper leaves. Second, the path connecting the two leaves in R labeled by a and c contains at least one pendant proper leaf.

Preparing cherries. Given a rooted (nonbinary) phylogenetic $\mathcal{X}$-tree R, a forest $\mathcal{F}$ on some taxa set $\mathcal{X}'$ such that $\overline{\mathcal{F}}$ is a forest for R as well as a cherry $\{a, c\}$, then, $\{a, c\}$ *is prepared* as follows. If $\{a, c\}$ is not a pseudo cherry for R and $\mathcal{F}$, nothing is done. Else, the following two steps are conducted. First, each pendant proper leaf $\{\ell_1, \ldots, \ell_n\}$ in R lying on the path connecting the two leaves labeled by a and c is removed. Second, the two nodes in $\mathcal{F}$ labeled by a and c are cut. Notice that, after preparing $\{a, c\}$, the node $\mathrm{LCA}_R(\{a, c\})$ in R is the parent of the two leaves labeled by taxon a and c and each component in $\mathcal{F}$ referring to $\{\ell_1, \ldots, \ell_n\}$ only consists of a single isolated node (cf. Fig. 9).

Now, similar to the original algorithm, the algorithm allMulMAFs* is called by the same four parameters R, $\mathcal{F}$, M and k. Given two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2, R is initialized by T_1, $\mathcal{F}$ by $\{T_2\}$

and M by $\emptyset$. Depending on the size of $\mathcal{F}$ (cf. Case 1a–c) and the choice of the next cherry that is selected from R (cf. Case 2), the following steps are performed.

Case 1a. If F contains more than k components, the computational path is aborted immediately and an empty set is returned.

Case 1b. If R only consists of a single leaf, each F_i in $\mathcal{F}$ is expanded with the help of M, and, finally, returned.

Case 1c. If there exists a specific leaf ℓ in R that refers to an isolated node in F, this leaf ℓ is removed from R resulting in R'. Next, the algorithm branches into a new path by recursively calling the algorithm with R', $\mathcal{F}$, k, and M' corresponding to M updated by $M[\mathcal{L}(\ell)] \leftarrow \bot$.
Otherwise, if such a leaf ℓ does not exist continue with Case 2.

Case 2. Select a subtree in R in which each pair of taxa either represents a cherry or a pseudo cherry. Now, *for each* (pseudo) cherry $\{a, c\}$, depending on the location of the leaves referring to a and c in $\mathcal{F}$, one of the following three cases is performed. In a first step, however, the chosen cherry $\{a, c\}$ is prepared as described above. Moreover, M is updated by $M[\mathcal{X}_i] \leftarrow \bot$, where $\mathcal{X}_i$ denotes the taxa set of each proper leaf that has been cut during the preceding preparation step.

Case 2a. If $\{a, c\}$ is a common cherry, $\{a, c\}$ is processed as described in Case 2 corresponding to the original algorithm allMulMAFs.

Case 2b. If $a \not\sim_{\mathcal{F}} c$ holds, and, thus, the leaves referring to a and c in $\mathcal{F}$ are located in two different components, $\{a, c\}$ is processed as described in Case 3a corresponding to the original algorithm allMulMAFs.

Case 2c. If $a \sim_{\mathcal{F}} c$ holds, and, thus, the leaves referring to a and c in $\mathcal{F}$ can be found in the same component F_i, $\{a, c\}$ is processed as described in Case 3b corresponding to the original algorithm allMulMAFs.

Notice that there are two main differences between the algorithm allMulMAFs* and the original algorithm allMulMAFs. First, a computational path corresponding to allMulMAFs* can additionally process pseudo cherries. Second, if during a recursive call R contains a common cherry $\{a, c\}$ as well as a contradicting cherry $\{a, b\}$, allMulMAFs* additionally branches into a computational path processing $\{a, b\}$. In the following, we will call such a cherry $\{a, b\}$ a *needless cherry* as we will show later that it can be neglected for the computation of maximum agreement forests.

3.3.2 The algorithm ProcessCherries Lastly, we present a simplified version of the algorithm allMulMAFs* — called ProcessCherries — mimicking one of its computational by a *cherry list* $\bigwedge = (\wedge_1, \wedge_2, \ldots, \wedge_n)$, in which each of its elements $\wedge_i$ denotes a *cherry action*. Such a cherry action $\wedge_i = (\{a, c\}, \phi_i)$ is a tuple that contains a (pseudo) cherry $\{a, c\}$ of the corresponding rooted phylogenetic $\mathcal{X}$-tree R_i and the forest $\mathcal{F}_i$ as well as a variable $\phi_i \in \{\cup_{ac}, \nmid_a, \nmid_c, \cap_{ac}\}$ denoting the way $\{a, c\}$ is processed in iteration i. More precisely,

- $\cup_{ac}$ refers to contracting the cherry $\{a, c\}$ following Case 2a of the algorithm allMulMAFs*.
- $\nmid_a$ and $\nmid_c$ refers to cutting taxon a and c, respectively, of the cherry $\{a, c\}$ following Case 2b of the algorithm allMulMAFs*.
- $\cap_{ac}$ refers to cutting each pendant subtree connecting taxon a and taxon c in $\mathcal{F}_i$ following Case 2c of the algorithm allMulMAFs*.

Now, given a cherry list $\bigwedge$, we say that $\bigwedge$ is a *cherry list for* T_1 *and* T_2, if in each iteration i the cherry $\{a, c\}$ of $\wedge_i = (\{a, c\}, \phi_i)$ is either contained in R_i or $\{a, c\}$ is a pseudo cherry for R_i and F_i. Moreover, after having prepared the cherry $\{a, c\}$, one of the following two conditions has to be satisfied.

Algorithm 1: ProcessCherries$(R, \mathcal{F}, (\wedge_1, \ldots, \wedge_n))$

```
M ← ∅;
for i = 1, ..., n do
    ({a, c}, φ_i) ← ∧_i;
    if {a, c} is a cherry of R or a pseudo cherry for R and F then
        if {a, c} is a pseudo cherry for R and F then
            (R, F, M) ← prepare pseudo cherry {a, c};
        if {a, c} is a common cherry of R and F and φ_i == ∪_ac then
            (R, F, M) ← contract cherry {a, c};
        else if {a, c} is a contradicting cherry of R and F and φ_i ==↾_a then
            e_a ← in-edge of node labeled by a in F;
            (R, F, M) ← cut edge e_a in F;
        else if {a, c} is a contradicting cherry of R and F and φ_i ==↾_c then
            e_c ← in-edge of node labeled by c in F;
            (R, F, M) ← cut edge e_c in F;
        else if {a, c} is a contradicting cherry of R and F then
            F ← F[a ∼ c];
            E_B ← set of pendant edges for a and c in F;
            (R, F, M) ← cut each edge in E_B;
        else
            return (∅);
        (R, F, M) ← from R remove each leaf referring to an isolated node in F;
    else
        return (∅);
F ← expand F as prescribed in M;
return (F);
```

– Either $\{a, c\}$ is a common cherry of R_i and F_i and $\phi_i = \cup_{ac}$,
– or $\{a, c\}$ is a contradicting cherry of R_i and F_i.

Notice that this is the case, if and only if calling ProcessCherries$(T_1, \{T_2\}, \bigwedge)$ does not return the empty set (cf. Alg. 1).

3.3.3 Proof of Correctness In this section, we will establish the correctness of the algorithm allMulMAFs by establishing the following theorem.

Theorem 1. *Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, by calling*

$$\text{allMulMAFs}(T_1, \{T_2\}, \emptyset, k)$$

all relevant maximum agreement forests for T_1 and T_2 are calculated, if and only if $k \geq h(T_1, T_2)$.

Proof. The proof of Theorem 1 is established in several substeps. First, given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, we will show that a binary resolution of each maximum agreement forest for T_1 and T_2 can be computed by applying the algorithm allMAFs^1 to a binary resolution of T_1 and T_2, where, as already mentioned, allMAFs^1 denotes a modification of the algorithm allMAAFs^1 omitting the acyclic check. Next, we will show that for an agreement forest $\mathcal{F}$ calculated by allMulMAFs there does not exist en edge e such that $\mathcal{F} \ominus \{e\}$ is still an agreement forest for T_1 and T_2, which directly implies that $\mathcal{F}$ is relevant. Moreover, we will show that, if a cherry list $\bigwedge$ for two binary resolutions $\hat{T}_1$ and $\hat{T}_2$ of T_1 and T_2, respectively, computes a maximum agreement forest $\hat{\mathcal{F}}$ for $\hat{T}_1$ and $\hat{T}_2$, $\bigwedge$ is mimicking a computational path of the algorithm allMulMAFs* calculating an agreement forest $\mathcal{F}$ for T_1 and T_2 such that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$. Lastly, we will show that each maximum agreement forest $\mathcal{F}$ computed by allMulMAFs* is also computed by allMulMAFs.

Lemma 1. *Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, for each relevant maximum agreement forest $\mathcal{F}$ for T_1 and T_2 of size k there exists a binary resolution $\hat{\mathcal{F}}$ of $\mathcal{F}$ that is calculated by*

$$\text{allMAFs}^1(\hat{T}_1, \hat{T}_2, \hat{T}_1, \hat{T}_2, k),$$

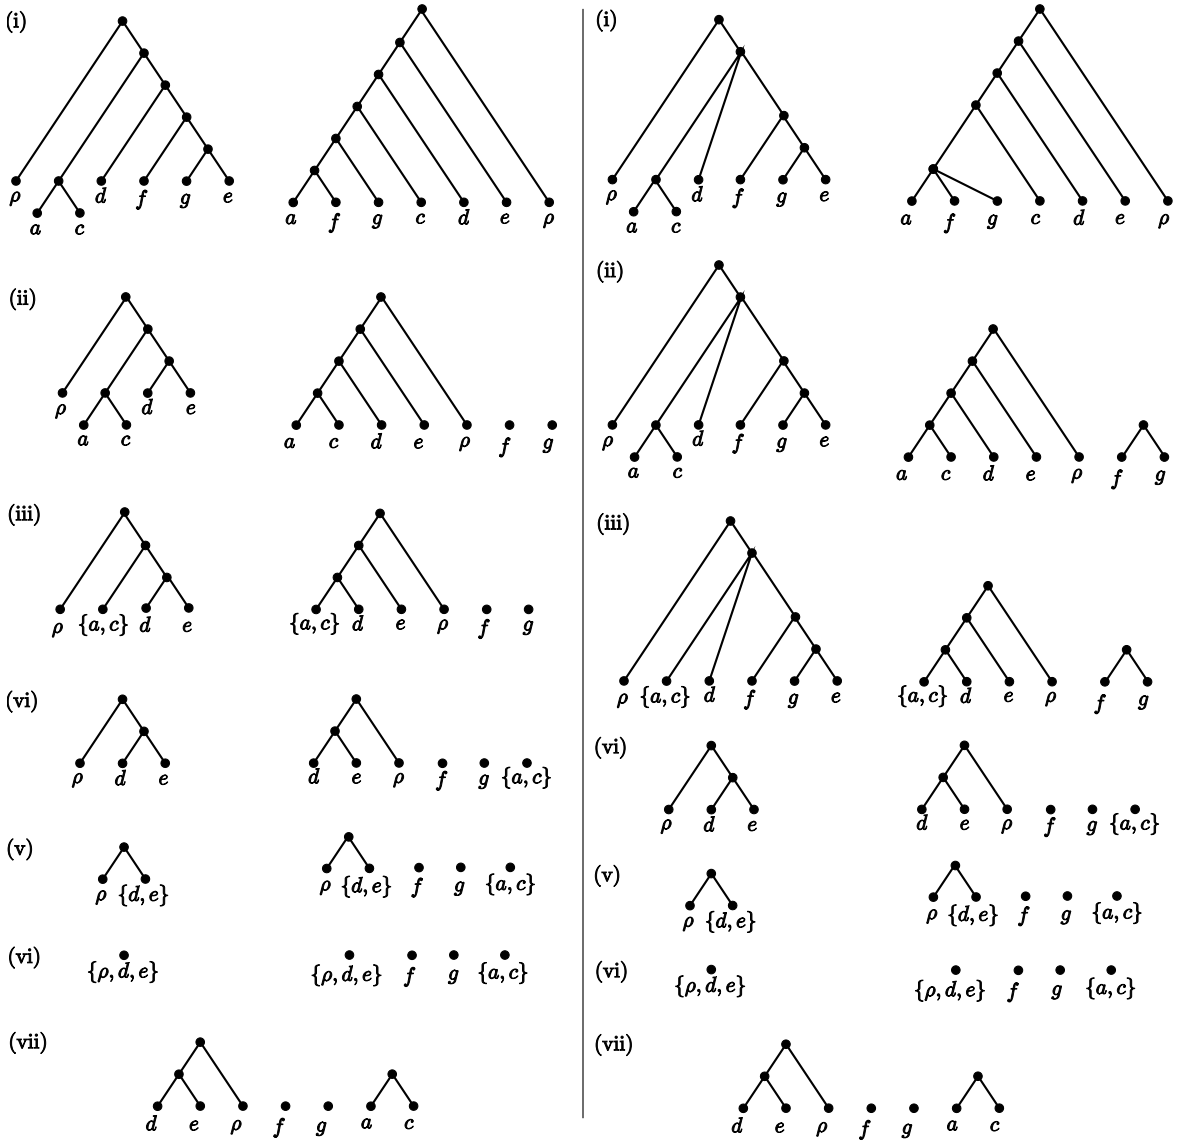

Fig. 10. Two examples of calling ProcessCherries for two binary and nonbinary trees according to the cherry actions $(\{a,c\},\cap_{ac}),(\{a,c\},\cup_{ac}),(\{d,e\},\cap_{de}),(\{d,e\},\cup_{de})$, and $(\{\rho,\{d,e\}\},\cup_{\rho de})$ conducted in sequential order. Notice that, as the two binary trees are binary resolutions of the two nonbinary trees, the resulting forest on the left hand side is a binary resolution of the resulting forest on the right hand side. Moreover, regarding Step (iii) on the right hand side, notice that the chosen cherry $\{d,e\}$ is a pseudo cherry and, thus, in Step (iv) the two components consisting of the single nodes labeled by f and g, respectively, arise from preparing $\{d,e\}$.

where $\hat{T}_1$ and $\hat{T}_2$ refers to binary resolutions of T_1 and T_2, respectively.

Proof. First notice that the algorithm allMAAFs[1] without conducting the acyclic check computes all maximum agreement forest for two rooted binary phylogenetic $\mathcal{X}$-trees, which is a direct consequence from [4][Lemma 2]. Moreover, given a relevant maximum agreement forest $\mathcal{F}$ for T_1 and T_2, a binary resolution $\hat{\mathcal{F}}$ of $\mathcal{F}$ is automatically a maximum agreement forest corresponding to $T_1(\hat{\mathcal{F}})$ and $T_2(\hat{\mathcal{F}})$. This is, in particular, the case, since just by definition each component $\hat{F}$ in $\hat{\mathcal{F}}$ is a subtree of $T_1(\hat{\mathcal{F}})$ and $T_2(\hat{\mathcal{F}})$ and, as in $\mathcal{F}$ all components are edge disjoint subtrees in T_1 and T_2, this has to hold for each of its binary resolution as well. Furthermore, $\hat{\mathcal{F}}$ has to be of minimum cardinality, since, otherwise, $\mathcal{F}$ would not be a maximum agreement forest for T_1 and T_2. Consequently, by applying the algorithm allMAFs[1] to both trees $T_1(\hat{\mathcal{F}})$ and $T_2(\hat{\mathcal{F}})$ the maximum agreement forest $\hat{\mathcal{F}}$ is calculated if $k \geq |\mathcal{F}|$, which, finally, establishes the proof of Lemma 1.

In the following, we will show that a cherry list $\bigwedge$ for two binary resolutions of two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2 is also mimicking a computational path of allMulMAFs* applied to T_1 and T_2.

Lemma 2. *Let $\hat{T}_1$ and $\hat{T}_2$ be two binary resolutions of two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, respectively. Moreover, let $\bigwedge$ be a cherry list for $\hat{T}_1$ and $\hat{T}_2$. Then, $\bigwedge$ is automatically a cherry list for T_1 and T_2.*

Proof. Lemma 2 obviously holds, if the cherry list $\bigwedge$ only exists of cherry actions $\wedge_i = (\{a, c\}, \phi_i)$ with $\phi_i \in \{\cup_{ac}, \restriction_a, \restriction_c, \}$. This is, in particular, the case because these cherry actions only affect those nodes whose corresponding subtree has been fully contracted so far. When processing a cherry action $\wedge_i = (\{a, c\}, \phi_i)$ with $\phi_i = \cap_{ac}$, however, two slightly different forests $\hat{\mathcal{F}}_{i+1}$ and $\mathcal{F}_{i+1}$ can arise. More precisely, this is the case, if there is a multifurcating node x lying on the path P_{ac} connecting taxon a and c in $\mathcal{F}_i$ providing a set $V_x = \{v_0, v_1, v_2, \ldots, v_n\}$ of at least 3 children, where v_0 denotes the node which is also part of P and, if, additionally, R_i contains two nodes d and e whose path connecting d and e contains a set of pendant subtrees each corresponding to $R_i(v_i)$, with $i > 0$ and $v_i \in V_x$.

Now, for simplicity, we assume that there is only one such multifurcating node x of out-degree three so that $V_x = \{v_0, v_1, v_2\}$. In this case, as $\hat{\mathcal{F}}_i$ only consists of binary trees, by processing $\wedge_i = (\{a, c\}, \cap_{ac})$ two components $\hat{F}_\alpha$ and $\hat{F}_\beta$ rooted at v_1 and v_2, respectively, are added to $\hat{\mathcal{F}}_{i+1}$ whereas to $\mathcal{F}_{i+1}$ only one component F_γ is added whose root contains two children corresponding to v_1 and v_2 (cf. Fig. 11).

Now, let $\wedge_j$ be a cherry action in $\bigwedge$ with $j > i$ in which both components $\hat{F}_\alpha$ and $\hat{F}_\beta$ have been fully contracted so far. As a consequence, since the components $\hat{F}_\alpha$ and $\hat{F}_\beta$ only consist of a single taxon, which has been removed from $\hat{R}_{j+1}$ (cf. Alg. 1, Line 13), the two taxa d and e are now cherries in $\hat{R}_{j+1}$ which is, however, not the case in R_{j+1} because F_γ still contains the two nodes v_1 and v_2 (cf. Fig. 11). Nevertheless, since in F_γ the node v_1 and v_2 are leaves directly attached to the root, the cherry $\{d, e\}$ is a pseudo cherry in R_{j+1} and, thus, an upcoming cherry action $\wedge_k$ containing $\{d, e\}$ represents a pseudo cherry for R_k and $\mathcal{F}_k$ in this case.

As a consequence, each cherry of $\hat{R}_i$ and $\hat{\mathcal{F}}_i$ corresponding to a cherry action $\wedge_i$ in $\bigwedge$ is either a cherry or a pseudo cherry of R_i and $\mathcal{F}_i$ and, thus, Lemma 2 is established.

Next, we will show that by expanding a forest $\mathcal{F}'$ on $\mathcal{X}$ as prescribed in M derived from calling allMulMAFs, the resulting forest is automatically an agreement forest for both input trees T_1 and T_2.

Lemma 3. *Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, let $\mathcal{F}$ be a forest on $\mathcal{X}$ that has been expanded as prescribed in M after allMulMAFs*$(T_1, \{T_2\}, \emptyset, k)$ has been called. Then, $\mathcal{F}$ is an agreement forest for T_1 and T_2.*

Proof. Since, obviously, $\mathcal{F}$ is a partition of $\mathcal{X}$, it suffices to consider each of the following cases describing a putative scenario leading to a forest that is not an agreement forest for both input trees T_1 and T_2, because either the refinement property or the node disjoint property in terms of T_1 or T_2 is not fulfilled. We will show, however, that during the execution of our algorithm allMulMAFs* each of those scenarios can be excluded.

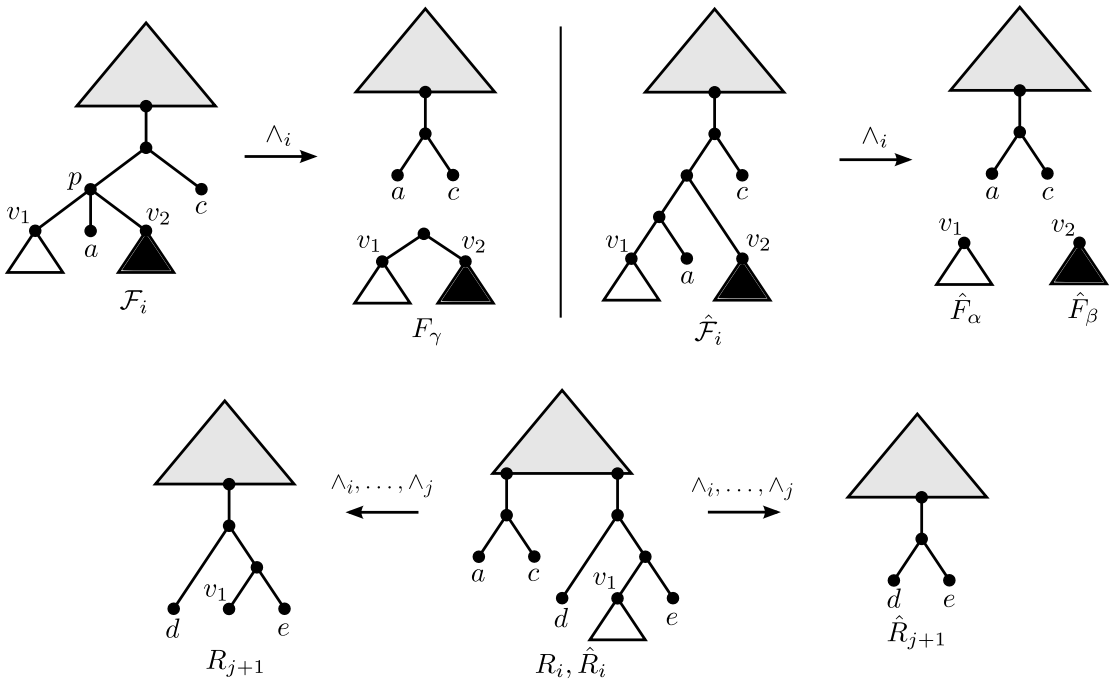

Fig. 11. An illustration of the scenario described in the proof of Lemma 2.

Case 1. Assume there exists a component F_i in $\mathcal{F}$ such that F_i is not a refinement of $T_2|_{\mathcal{X}_i}$, where $\mathcal{X}_i$ denotes the taxa set of F_i. As $\mathcal{F}'$ has been derived from T_2 by cutting and contracting its edges, this automatically implies that a cherry has been expanded as prescribed in M in respect of $\top$ instead of $\bot$. However, in M a cherry is only then mapped to $\top$, if and only if, during the i-th recursive call, it is a common cherry of R_i and $\mathcal{F}_i$ and if both parents corresponding to the cherry in R_i and F_i are multifurcating nodes (cf. Case 2a of allMulMAFs*). Moreover, such a cherry is immediately mapped back to $\bot$, if either the cherry itself or all its siblings have been cut (cf. Case 1c,2c of allMulMAFs*). Thus, such a component F_i cannot exist in $\mathcal{F}$.

Case 2. Assume there exists a component F_i in $\mathcal{F}$ such that F_i is not a refinement of $T_1|_{\mathcal{X}_i}$, where $\mathcal{X}_i$ denotes the taxa set of F_i. This automatically implies that either a cherry has been expanded as prescribed in M in respect of $\top$ instead of $\bot$ or, during the i-th recursive call, a cherry was not a common cherry of R_i and F_i. Similar to Case 1, the first scenario can be excluded. Moreover, the latter scenario cannot take place either, since, in order to reduce F_i to a single node, this common cherry must have been contracted which can only take place, if it was a common cherry of R_i and $\mathcal{F}_i$. Thus, such a component F_i cannot exist in $\mathcal{F}$.

Case 3. Assume there exist two components F_i and F_j in $\mathcal{F}$, with taxa set $\mathcal{X}_i$ and $\mathcal{X}_j$, respectively, such that $T_2(\mathcal{X}_i)$ and $T_2(\mathcal{X}_j)$ are not edge disjoint in T_2. As F_i and F_j must be a refinement of $T_2|_{\mathcal{X}_i}$ and $T_2|_{\mathcal{X}_j}$, respectively, (cf. Case 1) and both components have been derived from T_2 by cutting some of its edges, only one of both components can be part of $\mathcal{F}$. As a direct consequence, such two components cannot exist in $\mathcal{F}$.

Case 4. Assume there exist two components F_i and F_j in $\mathcal{F}$ such that $T_1(\mathcal{X}_i)$ and $T_1(\mathcal{X}_j)$ are not edge disjoint in T_1. As shown in Case 2, F_i and F_j must be a refinement of $T_1|_{\mathcal{X}_i}$ and $T_1|_{\mathcal{X}_j}$, respectively, and, thus, in order to obtain F_i and F_j the following steps must be performed during the execution of allMulMAFs*. Let E_i be an edge set that is only contained in $T_1(\mathcal{X}_i)$ and not in $T_1(\mathcal{X}_j)$. In order to obtain F_i, some of those edges in E_j must be cut, whereas, in order to obtain F_j, all of them must be contracted which leads

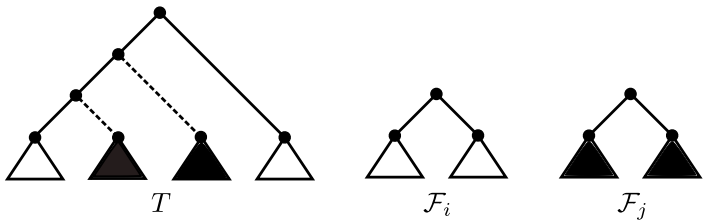

Fig. 12. An illustration of the scenario described in Case 4 corresponding to the proof of Lemma 3. In order to obtain F_i from T, both dashed edges have to be cut whereas, in order to obtain F_j, these two edges have to be contracted which is a contradiction.

to a contradiction (cf. Fig. 12). Thus, such two components cannot exist in $\mathcal{F}$.

Finally, by combining all four cases Lemma 3 is established.

Moreover, in the following, we will show that each agreement forest $\mathcal{F}$ that is reported by allMulMAFs* is relevant which means that $\mathcal{F}$ does not contain an edge e such that $\mathcal{F} \ominus \{e\}$ is still an agreement forest for both input trees.

Lemma 4. *Let T_1 and T_2 be two rooted (nonbinary) phylogenetic trees, then, each agreement forest that is reported by applying* allMulMAFs* *to T_1 and T_2 is relevant.*

Proof. Just by definition, given two rooted (nonbinary) phylogenetic trees T_1 and T_2, an agreement forest $\mathcal{F}$ for T_1 and T_2 that is not relevant has to contain an edge e such that $\mathcal{F} \ominus \{e\}$ is still an agreement forest for T_1 and T_2. Such an edge e, however, can only arise, if a cherry of a multifurcating node is expanded in respect of $\bot$ instead of $\top$. Initially, such a cherry must have been set to $\top$ because during the i-th recursive call both corresponding parents in R_i and $\mathcal{F}_i$, respectively, must have been multifurcating nodes. The only scenario setting $\top$ to $\bot$ would arise, if the cherry itself or all its siblings are cut during subsequent recursive calls. In this case, however, this cherry has to be set to $\bot$, since, otherwise the resulting forest would not be an agreement forest for T_1 and T_2. Thus, such an edge e cannot exist and, consequently, Lemma 4 is established.

Since both algorithms allMulMAFs* and allMulMAFs process common cherries and contradicting cherries in the same way, Lemma 4, obviously, has to hold for allMulMAFs as well.

Corollary 1. *Let T_1 and T_2 be two rooted (nonbinary) phylogenetic trees, then, each agreement forest that is reported by applying* allMulMAFs *to T_1 and T_2 is relevant.*

Now, let T_1 and T_2 be two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees. By the following Lemma 5, we will show that for each maximum binary agreement forest $\hat{\mathcal{F}}$, which can be computed by applying ProcessCherries to a cherry list $\bigwedge$ for two binary resolutions of T_1 and T_2, by calling ProcessCherries$(T_1, T_2, \bigwedge)$ a forest $\mathcal{F}$ is computed such that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$.

Lemma 5. *Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, let $\hat{T}_1$ and $\hat{T}_2$ be two binary resolutions of T_1 and T_2, respectively. Moreover, let $\hat{\mathcal{F}}$ be an agreement forest for $\hat{T}_1$ and $\hat{T}_2$ obtained from calling* ProcessCherries$(\hat{T}_1, \{\hat{T}_2\}, \bigwedge)$*, where $\bigwedge$ denotes a cherry list for $\hat{T}_1$ and $\hat{T}_2$. Then, a relevant agreement forest $\mathcal{F}$ is calculated by calling* ProcessCherries$(T_1, \{T_2\}, \bigwedge)$ *such that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$.*

Proof. We will first show by induction a slightly modified version of Lemma 5.

Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2. Let $\hat{T}_1$ and $\hat{T}_2$ be two binary resolutions of T_1 and T_2, respectively, and let $\hat{\mathcal{F}}_i$ and $\mathcal{F}_i$ be each forest corresponding to iteration i while executing

$$\text{ProcessCherries}(\hat{T}_1, \{\hat{T}_2\}, \bigwedge) \text{ and } \text{ProcessCherries}(T_1, \{T_2\}, \bigwedge),$$

respectively, where $\bigwedge = (\wedge_1, \wedge_2, \ldots, \wedge_n)$ is a cherry list for $\hat{T}_1$ and $\hat{T}_2$. Then, $\hat{\mathcal{F}}_i$ is a called *pseudo binary resolution* of $\mathcal{F}_i$, which is defined as follows. Given two forests $\hat{\mathcal{F}}$ and $\mathcal{F}$ for a phylogenetic $\mathcal{X}$-tree, we say that $\hat{\mathcal{F}}$ is a pseudo binary resolution of $\mathcal{F}$, if for each component $\hat{F}$ in $\hat{\mathcal{F}}$ there exists a component F in $\mathcal{F}$ such that one of the two following properties hold.

(i) $\hat{F}$ is a binary resolution of F.
(ii) $\hat{F}$ is a binary resolution of $F(v)$, where v is a child of the root of F.

The following proof is established by an induction on i denoting the position of a cherry action in $\bigwedge = (\wedge_1, \wedge_2, \ldots, \wedge_n)$.

Base case. At the beginning, F_1 only consists of $\hat{T}_2$, which is a binary resolution of T_2. Thus, the assumption obviously holds for $i = 1$.

Inductive step. Depending on the cherry action $\wedge_i = (\{a, c\}, E)$, the forest $\hat{\mathcal{F}}_{i+1}$ can be obtained from $\hat{\mathcal{F}}_i$ in the following ways.

(i) If $\{a, c\}$ is a pseudo cherry, a set of nodes V' that is attach to the root of a component in $\mathcal{F}_i$ is cut. Since $\bigwedge$ is a cherry list for $\hat{T}_1$ and $\hat{T}_2$ and, thus, $\{a, c\}$ is a cherry in $\hat{R}_i$, each of node in V' already refers to components in $\hat{\mathcal{F}}_i$ all consisting only of isolated nodes. Thus, after cutting the in-edge of each node in V', $\hat{\mathcal{F}}_i$ is still a pseudo binary resolution of $\mathcal{F}_i$.
(ii) If $\{a, c\}$ is a common cherry and, thus $\phi_i = \cup_{ac}$, in both forests $\hat{\mathcal{F}}_i$ and $\mathcal{F}_i$ the two taxa a and c are contracted. Consequently, since $\hat{\mathcal{F}}_i$ is a pseudo binary resolution of $\mathcal{F}_i$, this directly implies that $\hat{\mathcal{F}}_{i+1}$ is a pseudo binary resolution of $\mathcal{F}_{i+1}$ as well.
(iii) If $\{a, c\}$ is a contradicting cherry and $\phi_i = \not|_a$ (or $\phi_i = \not|_c$), then, in both forests $\hat{\mathcal{F}}_i$ and $\mathcal{F}_i$ the node labeled by taxon a (or taxon c) is cut. Again, no matter if $a \sim c$ or $a \not\sim c$ holds, since $\hat{\mathcal{F}}_i$ is a pseudo binary resolution of $\mathcal{F}_i$, this directly implies that $\hat{\mathcal{F}}_{i+1}$ is a pseudo binary resolution of $\mathcal{F}_{i+1}$ as well.
(iv) If $\{a, c\}$ is a contradicting cherry and $\phi_i = \cap_{ac}$, in both forests $\hat{\mathcal{F}}_i$ and $\mathcal{F}_i[a \sim c]$ each pendant subtree lying on the path connecting both leaves labeled by a and c, respectively, is cut. Let $\hat{\mathcal{F}}'$ and $\mathcal{F}'$ be those component arising from cutting $\hat{\mathcal{F}}_i$ and $\mathcal{F}_i[a \sim c]$, respectively. Since $\hat{\mathcal{F}}_i$ is a binary resolution of $\mathcal{F}_i$, $|\mathcal{F}'| \geq |\hat{\mathcal{F}}'|$ holds which means, in particular, that each $\hat{F}'_i$ in $\hat{\mathcal{F}}'$ is either a binary resolution of F'_j or a binary resolution of $F'_j(v)$ in $\mathcal{F}'$, where v corresponds to a child whose parent is the root of F'_j. Thus, since $\hat{\mathcal{F}}_i$ is a pseudo binary resolution of $\mathcal{F}_i$, this directly implies that $\hat{\mathcal{F}}_{i+1}$ is a pseudo binary resolution of $\mathcal{F}_{i+1}$ as well.

Now, from the induction we can deduce that, independent from the cherry action $\wedge_i$, $\hat{\mathcal{F}}_i$ is always a pseudo binary resolution of $\mathcal{F}_i$. Moreover, let $\hat{\mathcal{F}}_{n+1}$ and $\mathcal{F}_{n+1}$ be the two forests obtained from $\hat{\mathcal{F}}_n$ and $\mathcal{F}_n$, respectively, by applying $\wedge_n$. Then, since $\hat{\mathcal{F}}$ is an agreement forest for $\hat{T}_1$ and $\hat{T}_2$, all components in $\hat{\mathcal{F}}_{n+1}$ only consist of single isolated nodes which directly implies that $\mathcal{F}_{n+1}$ does not contain any cherries. Furthermore, due to Lemma 3, by expanding $\mathcal{F}_{n+1}$ as prescribed in M a relevant agreement forest $\mathcal{F}$ arises such that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$ which completes the proof of Lemma 5.

In the following, we will show that Lemma 5 also holds for the original algorithm allMulMAFs.

Lemma 6. *Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, let $\hat{\mathcal{F}}$ be a binary maximum agreement forest for T_1 and T_2. Then, by calling*

$$\text{allMulMAFs}(T_1, \{T_2\}, \emptyset, k)$$

a relevant maximum agreement forests $\mathcal{F}$ for T_1 and T_2 is computed such that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$, if and only if $k \geq h(T_1, T_2)$.

Proof. Notice that, as proven in Lemma 5, Theorem 6 holds for the modified algorithm allMulMAFs*. Thus, in order to establish Lemma 6, we just have to show that the following two differences between both algorithms allMulMAFs* and allMulMAFs do not have an impact on the computation of maximum

agreement forests.

Needless cherries. First of all, let $\bigwedge$ be a cherry list for T_1 and T_2 mimicking a computational path of allMulMAFs* resulting in an agreement forest $\mathcal{F}$ for T_1 and T_2. Moreover, let $\bigwedge$ contain a cherry action $\wedge_i = (\{a, b\}, \phi_i)$ in which $\{a, b\}$ is a contradicting cherry of R_i and $\mathcal{F}_i$. Now, if R_i contains a taxon c such that $\{a, c\}$ is a common cherry, we call $\{a, b\}$ a *needless cherry*. Notice that a computational path corresponding to the original algorithm allMulMAFs does not consider needless cherries as it always prefers common cherries to contradicting cherries. In the following, however, we will show that for the computation of maximum agreement forests each computational path processing needless cherries can be neglected.

Let $\mathcal{F}$ be an agreement forest resulting from a computational path of allMulMAFs* in which, instead of processing a common cherry $\{a, c\}$, a needless cherry $\{a, b\}$ is processed by the cherry action $\wedge_i = (\{a, b\}, \nmid_a)$. This implies that $\mathcal{F}$ contains a component F_a corresponding to the expanded taxon a, which has been cut during the i-th iteration. Moreover, let F_c be the component in $\mathcal{F}$ containing the node v_c corresponding to taxon c in $\mathcal{F}_i$. Since $\{a, c\}$ has been a common cherry in iteration i, by attaching F_a back to the in-edge of v_c an agreement forest of size $k - 1$ arises and, thus, $\mathcal{F}$ cannot be a maximum agreement forest. Consequently, from cutting instead of contracting common cherries a maximum agreement forest cannot arise and, thus, for the computation of maximum agreement forests each computational path of allMulMAFs* processing needless cherries can be neglected.

Notice that, in this case, the cherry action $\wedge_i = (\{a, b\}, \nmid_b)$ would be also not be considered. However, after having contracted the common cherry $\{a, c\}$, b could still be cut selecting a cherry action involving one of its siblings.

Pseudo cherries. Furthermore, in contrast to the modified algorithm allMulMAFs*, a computational path corresponding to the original algorithm allMulMAFs does not consider pseudo cherries. In the following, however, we will show that for an agreement forest $\mathcal{F}$ resulting from a computational path processing pseudo cherries, there exists a different computational path calculating $\mathcal{F}$ without considering any pseudo cherries.

Let $\bigwedge$ be a cherry list for T_1 and T_2 mimicking a computational path of allMulMAFs* resulting in an agreement forest $\mathcal{F}$ and let $\wedge_i$ be a cherry action whose corresponding cherry $\{a, c\}$ is a pseudo cherry of R_i and $\mathcal{F}_i$. Moreover, let be $(b_1, b_2, \ldots, b_k)$ and $(b_{k+1}, b_2, \ldots, b_n)$ be those taxa corresponding to each pendant node lying on the path connecting a and $\mathrm{LCA}_{R_i}(\{a, c\})$ as well as c and $\mathrm{LCA}_{R_i}(\{a, c\})$, respectively. Then, we can replace $\wedge_i = (\{a, c\}, \phi_i)$ through the sequence of cherry actions

$$(\{a, b_1\}, \nmid_{b_1}), \ldots, (\{a, b_k\}, \nmid_{b_k}), (\{c, b_{k+1}\}, \nmid_{b_{k+1}}), \ldots, (\{c, b_n\}, \nmid_{b_n}), (\{a, c\}, \phi_i)$$

neither containing pseudo cherries nor needles cherries such that the agreement forest $\mathcal{F}$ is still computed. This means, in particular, that each tree operation that is conducted for preparing a pseudo cherry can be also realized by a sequence of cherry actions neither containing needless cherries nor pseudo cherries.

As shown above, for a relevant maximum agreement forest $\mathcal{F}$ our modified algorithm allMulMAFs* always contains a computational path calculating $\mathcal{F}$ by neither taking needless cherries nor pseudo cherries into account. Thus, each relevant maximum agreement forest for T_1 and T_2 that is calculated by allMulMAFs* is also calculated by allMulMAFs and, as a direct consequence, Lemma 6 is established.

Now, in a last step, we can finish the proof of Theorem 1. Let $\hat{T}_1$ and $\hat{T}_2$ be two binary resolutions of two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2, respectively, and let $\hat{\mathcal{F}}$ be an agreement forest for $\hat{T}_1$ and $\hat{T}_2$. Then, by combining Corollary 1 and Lemma 6 we can deduce that the algorithm allMulMAFs computes a relevant maximum agreement forest $\mathcal{F}$ for T_1 and T_2 such that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$. This automatically implies, that our algorithm calculates all relevant maximum agreement forests for T_1 and T_2 and, thus, Theorem 1 is finally established.

3.4 Runtime of allMulMAFs

In this section, we discuss the theoretical worst-case runtime of the algorithm allMulMAFs in detail.

Theorem 2. *Let T_1 and T_2 be two rooted phylogenetic $\mathcal{X}$-trees and $\mathcal{F}$ be a relevant maximum agreement forest for T_1 and T_2 containing k components. The theoretical worst-case runtime of the algorithm* allMulMAFs *applied to T_1 and T_2 is $O(3^{|\mathcal{X}|+k}|\mathcal{X}|)$.*

Proof. Let $\mathcal{F} = \{F_\rho, F_1, F_2, \ldots, F_{k-1}\}$ be an agreement forest for T_1 and T_2 of size k. To obtain $\mathcal{F}$ from T_2, obviously $k-1$ edge cuttings are necessary. Moreover, in order to reduce the size of the leaf set $\mathcal{X}$ of R to 1, to each component F_i in $\mathcal{F}$ we have to apply exactly $|\mathcal{L}(F_i)| - 1$ cherry contractions. Consequently, at most $|\mathcal{X}|$ cherry contractions have to be performed in total. Thus, our algorithm has to perform at most $O(|\mathcal{X}| + k)$ recursive calls for the computation of $\mathcal{F}$. Now, as one of these recursive calls can at least branch into 3 further recursive calls, $O(3^{|\mathcal{X}|+k})$ is an upper bound for the total number of recursive calls that are performed throughout the whole algorithm. Moreover, each case that is conducted during a recursive (cf. Sec. 3.2) can be done in $O(|\mathcal{X}|)$ time and, thus, the theoretical worst-case runtime of the algorithm can be estimated with $O(3^{|\mathcal{X}|+k}|\mathcal{X}|)$.

3.5 Conclusion

In this section, we have presented the algorithm allMulMAFs calculating all relevant maximum agreement forests for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees. Therefor, we have established a detailed formal proof showing the correctness of the algorithm which is based on both previously presented algorithms allMAAFs and allMAAFs[1]. In the next section, we will show how to further modify the algorithm allMulMAFs so that now all relevant maximum *acyclic* agreement forests are calculated.

4 The algorithm allMulMAAFs

In this section, we show how to extend the algorithm allMulMAFs, presented in Section 3.2, such that the reported agreement forests additional satisfy the acyclic constraint, which automatically implies that the extended algorithm will calculated all relevant maximum *acyclic* agreement forests for two rooted *nonbinary* phylogenetic $\mathcal{X}$-trees. As mentioned previously, the acyclic constraint plays an important role for the construction of hybridization networks as, for example, demonstrated by the algorithm HybridPhylogeny [6]. More specifically, this algorithms generates a hybridization network displaying two rooted *bifurcating* phylogenetic $\mathcal{X}$-trees from the components of an acyclic agreement forest of those two trees. Thus, we consider the computation of nonbinary maximum acyclic agreement forests as a first step to come up with minimum hybridization networks displaying the refinements of two rooted nonbinary phylogenetic $\mathcal{X}$-trees.

Broadly speaking, the algorithm allMulMAFs can be used to make progress towards an agreement forest for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2 as long as the set of components $\mathcal{F}$ does not satisfy all properties of an agreement forest. Once our algorithm has successfully computed a maximum agreement forest $\mathcal{F}$ for T_1 and T_2, we can apply a specific tool that is able to check, if we can refine $\mathcal{F}$ to a maximum acyclic agreement forest. Such a refinement of an agreement forest is done by cutting a minimum number of edges within its components such that each directed cycle of the underlying ancestor-descendant graph $AG(T_1, T_2, \mathcal{F})$ is dissolved.

Notice that this problem is closely related to the *directed feedback vertex set problem*. More specifically, given a directed graph G with node set V, a feedback vertex set V' is a subset of V containing at least one node of each directed cycle of G. This implies, by deleting each node of V' together with its adjacent edges, each directed cycle is automatically removed. Now, based on a directed graph, the directed feedback vertex set problem consists of minimizing the size of such a feedback vertex set.

4.1 Refining agreement forests

In this section we present a tool that enables the refinement of an agreement forest. We call this tool an *expanded ancestor-descendant graph*. Notice that this tool has been previously published under a different term as we state in the following.

Remark 3. The following concept of an *expanded ancestor-descendant graph* corresponds to the concept of an *expanded cycle graph* given in the work of Whidden *et al.* [15]. The latter concept, however, can be only applied to agreement forests corresponding to rooted *binary* phylogenetic $\mathcal{X}$-trees. Hence, we have adapted this concept such that it can be also applied to agreement forests corresponding to rooted *nonbinary* phylogenetic $\mathcal{X}$-trees. Notice that, adapting the concept of an expanded cycle graph to nonbinary agreement forests has been also examined in the master thesis of Li [12]. In this work, each step that is necessary to compute the hybridization number for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees is presented in detail by, additionally, discussing its correctness.

In the following, we give a short overview of how an expanded ancestor-descendant graph is defined and how this graph can be used to transform agreement forests into acyclic agreement forests.

Expanded ancestor-descendant graph. The tool that enables the refinement of an agreement forest $\mathcal{F}$ for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2 to an acyclic agreement forest is an *expanded ancestor-descendant graph* $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$. In contrast to the ancestor-descendant graph, each node of this graph corresponds to exactly one particular node of a component in $\mathcal{F}$. Thus, from such a graph one can directly figure out those edges of a component that have to be cut in order to remove a directed cycle (cf. Fig. 13).

Given two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2 and a nonbinary maximum acyclic agreement forest $\mathcal{F}$ for T_1 and T_2, the corresponding expanded ancestor-descendant graph $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ consists of the following nodes and edges. First of all, $\mathcal{F}$ is a subset of $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$, which means that the graph contains all nodes and edges corresponding to all components in $\mathcal{F}$. Moreover, $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ contains a set of *hybrid edges* each connecting two specific nodes each being part of two different components. More precisely, those edges are defined as follows.

Given a node v of a component F_j in $\mathcal{F}$, the function $\phi_i(v)$ refers to the lowest common ancestor in $T_i(\mathcal{F})$, with $i \in \{1, 2\}$, of each leaf that is labeled by a taxon contained in $\mathcal{L}(F_j(v))$. Notice that the node $\phi_i(\cdot)$ is well defined, which means there exists exactly one node in $T_i(\mathcal{F})$ to which $\phi_i(\cdot)$ applies. Equivalently, the function $\phi_i^{-1}(\cdot)$ maps nodes from T_i, with $i \in \{1, 2\}$, back to a component in $\mathcal{F}$. More precisely, let E_i, with $i \in \{1, 2\}$, be the set of edges consisting of all in-edges of all lowest common ancestors in T_i of the taxa set of each F_j in $\mathcal{F} \setminus \{F_\rho\}$. Then, the node $v \in T_i$ maps back to the node in $\mathcal{F}$ representing the lowest common ancestor of those taxa that can be reached from v by not using an edge in E_i. Notice, however, that this function is only defined for those nodes that are either labeled or are part of a path connecting two labeled nodes a and b such that $\phi_i^{-1}(a)$ and $\phi_i^{-1}(b)$ are contained in the same component F_j in $\mathcal{F}$. Similar to the binary case, since the graph is built for the trees $T_1(\mathcal{F})$ and $T_2(\mathcal{F})$ reflecting $\mathcal{F}$, the function $\phi_i^{-1}(\cdot)$ is well defined, which means that, if defined, there exists exactly one node in $\mathcal{F}$ to which $\phi_i^{-1}(\cdot)$ applies.

Now, based on the definitions of these two functions, $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ contains the following hybrid edges. Let w be a node in this graph corresponding to the root of a component F_j not equal to F_ρ. Moreover, for the tree $T_i(\mathcal{F})$ with $i \in \{1, 2\}$, let v' be the lowest ancestor of $\phi_i(w)$ such that $\phi_i^{-1}(v')$ is defined. In more detail, let $P_\phi = (v_1, \ldots, v_n)$ be those nodes lying on the path connecting the parent v_1 of v and the root v_n of T_1 such that v_j with $j \in [2 : n]$ is the parent of v_{j-1}. Then,

$$v' = \min_j \{v_j : v_j \in P_\phi \wedge \phi_i^{-1}(v_j) \text{ is defined}\}.$$

Based on v' and w, $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ contains a hybrid edge $(\phi_i^{-1}(v'), w)$. Notice that, if $\mathcal{F}$ contains k components, for each component except F_ρ two hybrid edges corresponding to T_1 and T_2 are inserted which are $2k - 2$ hybrid edges in total. Furthermore, the target node of a hybrid edge does always refer to a root node of a component F_j in $\mathcal{F}$ whereas the source node never does.

Exit nodes. Given two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2 as well as a nonbinary maximum acyclic forest $\mathcal{F}$ for T_1 and T_2, an *exit node of* $AG^{ex}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ is defined as follows. Let H_i be the set of hybrid edges in $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ resulting from T_i with $i \in \{1, 2\}$. Now, given a directed cycle in $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ running through the hybrid edges $E_H = \{h_0, \ldots, h_{n-1}\}$ in sequential order, then,

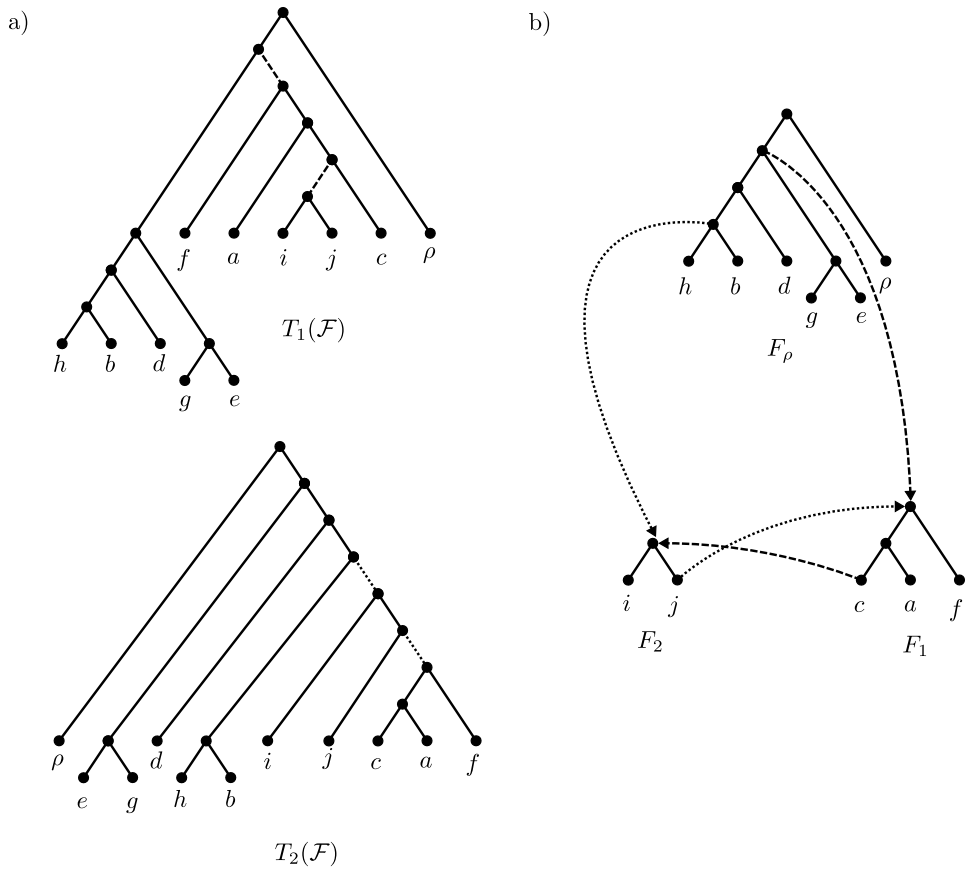

Fig. 13. (a) The same two trees $T_1(\mathcal{F})$ and $T_2(\mathcal{F})$ as depicted in Figure 3. Here, the set of dashed edges and the set of dotted edges refers to the in-edges of the nodes corresponding to the lowest common ancestors of $\mathcal{L}(F_1)$ and $\mathcal{L}(F_2)$. **(b)**The expanded ancestor-descendant graph $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ with $\mathcal{F} = \{F_\rho, F_1, F_2\}$. Dashed edges are hybrid edges resulting from $T_1(\mathcal{F})$ and dotted edges are hybrid edges resulting from $T_2(\mathcal{F})$. Here, for a better overview, the directions of edges corresponding to components of $\mathcal{F}$ are omitted. Notice that, by fixing the exit node corresponding to taxon j, all directed cycles are removed and a maximum acyclic agreement forest for T_1 and T_2 with size 4 arises.

the source node v_i of a hybrid edge $h_i = (v_i, w_i)$ in E_H is called an exit node, if h_i is contained in H_1 and h_j, with $j = (i-1) \mod n$, is contained in H_2 or vice versa.

Now, based on an expanded ancestor-descendant graph we can refine an agreement forest by *fixing its exit nodes*. An exit node v belonging to the component F_j is fixed by cutting each edge lying on the path connecting v with the node referring to the root node of F_j. Notice that by cutting k of those edges, the resulting agreement forest $\mathcal{F}'$ consists of $|\mathcal{F}| + k$ components.

4.2 The algorithm

We can easily turn the algorithm allMulMAFs into the algorithm allMulMAAFs by applying a post-processing step refining agreement forests. More precisely, given an agreement forest $\mathcal{F}$ for two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2, by applying the following refinement procedure only those relevant acyclic agreement forests are returned whose size is smaller than or equal to k.

(1) Compute two trees $T_1(\mathcal{F})$ and $T_2(\mathcal{F})$ reflecting $\mathcal{F}$.
(2) Build the expanded ancestor-descendant graph $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$.

(3) Compute the set of exit nodes V_H of $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$.
(4) For each exit node v in V_H turn $\mathcal{F}$ into $\mathcal{F}'$ by fixing v.
(5) For each agreement forest $\mathcal{F}'$ with $|\mathcal{F}'| \leq k$ continue with step 5a or 5b.
  (5a) If $\mathcal{F}'_i$ is acyclic, return $\mathcal{F}'$.
  (5b) Otherwise, if $\mathcal{F}'$ is *not* acyclic, repeat step 2–5 with $\mathcal{F}'$.

Based on these steps, by modifying Case 1b as follows, we can easily turn the algorithm allMulMAFs into an algorithm computing a set of maximum acyclic agreement forests.

Case 1b'. If R only consists of a single leaf, first each F_i in $\mathcal{F}$ is expanded as prescribed in M and then $\mathcal{F}$ is refined with the help of $AG^{\text{ex}}(T_1(\mathcal{F}), T_2(\mathcal{F}), \mathcal{F})$ into $\mathcal{F}'$. Finally, $\mathcal{F}'$ is returned.

This means that, each time before reporting an agreement forest $\mathcal{F}$, we first check, if we can refine $\mathcal{F}$ to an acyclic agreement forest $\mathcal{F}'$ of size smaller than or equal to k. If this is possible, we return $\mathcal{F}'$, else, we return the empty set.

4.3 Correctness of allMulMAAFs

In this section, we show that by applying the presented algorithm allMulMAAFs one can calculate all relevant maximum acyclic agreement forests for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees.

Theorem 3. *Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees, by calling*

$$\text{allMulMAAFs}(T_1, \{T_2\}, \emptyset, k)$$

all relevant maximum acyclic agreement forests for T_1 and T_2 are calculated, if and only if $k \geq h(T_1, T_2)$.

Proof. The correctness of the algorithm as stated in Theorem 3 directly depends on the following two Lemmas 7 and 8.

Lemma 7. *Let T_1 and T_2 be two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees and let $\mathcal{F}$ be a relevant maximum acyclic agreement forest T_1 and T_2. Then, a relevant agreement forest $\mathcal{F}'$ by calling* $\text{allMulMAFs}(T_1, T_2, \emptyset, h(T_1, T_2))$ *is calculated that can be turned into $\mathcal{F}$ by first resolving some of its multifurcating nodes and then by cutting some of its edges.*

Proof. As the first point holds for the algorithm allMAAFs^2 [4][Theorem 3], from Lemma 5 we can deduce that this has to hold for the algorithm allMulMAFs as well. More precisely, let T_1 and T_2 be two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees and let $\hat{\mathcal{F}}$ be a binary agreement forest for $T_1(\mathcal{F})$ and $T_2(\mathcal{F})$ that can be turned into a maximum acyclic agreement forest for $T_1(\mathcal{F})$ and $T_2(\mathcal{F})$ by cutting some of its edges. Then, due to Lemma 5, by calling $\text{allMAAFs}^2(T_1, T_2, \emptyset, h(T_1, T_2))$ a relevant acyclic agreement forest $\mathcal{F}$ for T_1 and T_2 is calculated such that $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$. Moreover, as $\hat{\mathcal{F}}$ can be turned into a maximum acyclic agreement forest by cutting some of its edges $\hat{E}$, $\mathcal{F}$ can be turned into a relevant maximum acyclic agreement forest as well by first resolving some of its nodes and then by cutting a certain edge set E with $|\hat{E}| = |E|$. More specifically, for each edge $\hat{e}$ in $\hat{E}$ there exists an edge e that can be obtained from $\mathcal{F}$ as follows. Let $\hat{e} = (\hat{v}, \hat{w})$ be an edge in $\hat{E}$ of a component $\hat{F}$ in $\hat{\mathcal{F}}$, then, as $\hat{\mathcal{F}}$ is a binary resolution of $\mathcal{F}$, $\mathcal{F}$ has to contain a component F with node w' such that $\mathcal{L}(\hat{F}(\hat{w})) \subseteq \mathcal{L}(F(w'))$. Now, e is the in-edge of a node w that can be obtained from resolving node w' such that $\mathcal{L}(F(\hat{w})) = \mathcal{L}(F(w))$.

Lemma 8. *Given two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees T_1 and T_2 as well as a relevant agreement forest $\mathcal{F}$ for T_1 and T_2, the refinement step resolves a minimum number of nodes and cuts a minimum number of edges such that $\mathcal{F}$ is turned into all relevant acyclic agreement forests of minimum size.*

Proof. Due to the following two observations that are both discussed in the master thesis of Li [12], the refinement procedure, which is based on fixing exit nodes as described above, leads to the computation of acyclic agreement forests.

Observation 1 *Let $\mathcal{F}$ be an agreement forest for two rooted phylogenetic $\mathcal{X}$-trees T_1 and T_2 and let $\mathcal{F}'$ be an agreement forest that is produced by fixing an exit node of $AG^{ex}(T_1, T_2, \mathcal{F})$. Then, the set of exit nodes corresponding to $AG^{ex}(T_1, T_2, \mathcal{F}')$ is a subset of the set of exit nodes corresponding to $AG^{ex}(T_1, T_2, \mathcal{F})$.*

Observation 2 *Given an agreement forest $\mathcal{F}$ for two rooted phylogenetic $\mathcal{X}$-trees, there exists an acyclic agreement forest $\mathcal{F}'$, if and only if there exists a set of exit nodes such that fixing theses nodes leads to the computation of $\mathcal{F}'$.*

A formal proof showing the correctness of these two observations can be looked up in the master thesis of Li [12]. More precisely, Observation 1 is a consequence of [12, Lemma 12 and 13], which ensures that by fixing an exit node one makes progress towards an acyclic agreement forest, and Observation 2 is a consequence of [12, Lemma 10], which ensures that it is possible to obtain all relevant maximum acyclic agreement forests from applying the refinement procedure. Notice that, as by fixing exit nodes a minimum number of nodes are resolved and a minimum number of edges are cut, each resulting maximum acyclic agreement forest is automatically relevant.

Now, from those two separate proofs each regarding two successive parts, namely the computation of specific relevant agreement forests followed by the refinement procedure establishing the acyclicity of each those forests, we can finally finish the proof of Theorem 3.

4.4 Runtime of allMulMAAFs

In this section, we discuss the runtime of the algorithm allMulMAAFs in detail.

Theorem 4. *Let T_1 and T_2 be two rooted phylogenetic $\mathcal{X}$-trees and $\mathcal{F}$ be a maximum agreement forest for T_1 and T_2 containing k components. The theoretical worst-case runtime of the algorithm* allMulMAAFs *applied to T_1 and T_2 is $O(3^{|\mathcal{X}|+k}4^k|\mathcal{X}|)$.*

Proof. As stated in Theorem 2, the algorithm has to conduct $O(3^{|\mathcal{X}|+k})$ recursive calls. Potentially, for each of those recursive calls we have to apply a refinement step whose theoretical worst-case runtime can be estimated as follows. First notice that the order of fixing exit nodes is irrelevant. Thus, in an expanded ancestor-descendant graph, corresponding to an agreement forest of size $k+1$ and, hence, containing $2k$ exit nodes, at most 2^{2k} different sets of potential exit nodes have to be considered. As the processing of such a set of potential exit nodes takes $O(|\mathcal{X}|)$ time, the theoretical worst-case runtime of the algorithm is $O(3^{|\mathcal{X}|+k}4^k|\mathcal{X}|)$.

In general, however, due to the following observation, the runtime of the refinement step is not a problem when computing maximum acyclic agreement forests of size k. Either the size k' of an agreement forest $\mathcal{F}$ is close to k and, thus, fixing an exit node immediately leads to an agreement forest of size larger than k (and, consequently, most of the sets of potential exit nodes have not to be considered in full extend). Otherwise, if the size k' of $\mathcal{F}$ is small and, thus, the gap between k' and k is large, the expanded ancestor-descendant graph is expected to contain no or at least only less cycles (and, consequently, there exist only few sets of potential exit nodes). Nevertheless, in the master thesis of Li [12] a method is presented that allows to half the number of exit nodes that have to be taken into account throughout the refinement of an agreement forest, so that by applying this modification the algorithm yields a theoretical worst-case runtime of $O(3^{|\mathcal{X}|+k}2^k k)$.

4.5 Robustness of our Implementation

In order to make the algorithm available for research, we added an implementation to our Java based software package Hybroscale providing a graphical user interface, which enables a user friendly interactive handling. Next, we conducted two specific test scenarios demonstrating the robustness of our implementation which means, in particular, that Hybroscale guarantees the computation of all relevant nonbinary acyclic agreement forests for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees. Each of those test scenarios was conducted on a particular synthetic dataset, which was generated as described below.

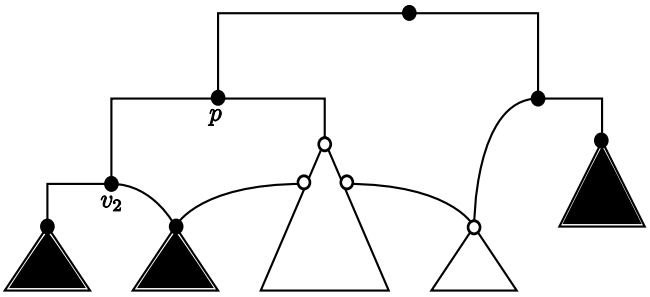

Fig. 14. An illustration of the cluster degree parameter. Given a cluster degree $c = 1$. When inserting an in-going edge e to node v_2 that is respecting c, each node that is marked white or is part of a white marked subnetwork forms a potential source node.

4.5.1 Synthetic dataset Our synthetic dataset consists of several tree sets each containing two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees. Each $\mathcal{X}$-tree is generated by ranging over all different combinations of four parameters, namely the number of leaves ℓ, an upper bound for the hybridization number k, the *cluster degree* c, and an additional parameter p. Each of both trees of a particular tree set corresponds to an embedded tree T of a particular network N only containing hybridization nodes of in-degree 2. With respect to the four different parameters such a tree T is computed as follows. First a random binary tree $\hat{T}$ containing ℓ leaves is computed. This is done, in particular, by randomly selecting two nodes u and v of a specific set V, which is initialized by creating ℓ nodes of both in- and out-degree 0. The two selected nodes u and v are then connected to a new node w. Finally, V is updated by replacing u and v by its parent node w. This is done until V only consists of one node corresponding to the root of $\hat{T}$. In a second step, k reticulation edges are inserted in $\hat{T}$ with respect to parameter c such that the resulting network N contains precisely k reticulation nodes of in-degree 2. Finally, after extracting a binary T' from N, based on parameter p, a certain percentage of its edges are contracted such that a nonbinary tree T is obtained from T'.

In this context, the *cluster degree* is an *ad hoc* concept influencing the computational complexity of a tree set similar to the concept of the *tangling degree* first presented in the work of Albrecht *et al.* [1] (cf. Fig. 14). When adding a reticulation edge e with target node v_2 and source node v_1, we say that e respects the cluster degree c, if v_1 cannot be reached from v_2 and there is a path of length less than or equal to c leading from v_2 to a certain node p such that v_1 can be reached from p. This means, in particular, that networks respecting a small cluster degree, in general, contain more minimum common clusters than networks respecting a large cluster degree and, thus, often provide a smaller computational complexity when applying a cluster reduction beforehand.

4.5.2 Comparison with other software First, we generated a synthetic dataset, as described above, containing tree pairs each consisting of two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees with parameters $\ell \in \{10, 25, 50\}$, $k \in \{5, 10, 15\}$, $c \in \{1, 3, 5\}$, and $p \in \{30\}$. More specifically, for all 81 combinations of the four parameters 30 tree sets were generated resulting in 810 tree sets in total. Next, based on this dataset, we compared the result of our implementation to the two software packages Dendroscope[1] [10] and TerminusEst[2] [13] so far being the only known available software packages computing exact hybridization numbers for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees.

Our simulation study pointed out that our implementation could always reproduce the hybridization numbers that were computed by both software packages Dendroscope and TerminusEst. Moreover, the number of maximum acyclic agreement forests computed by our algorithm was always larger than the number of networks that were reported by Dendroscope and TerminusEst. Notice that, regarding TerminusEst, this is not surprising as this program does only output one network. This fact, however, gives further

[1] `ab.inf.uni-tuebingen.de/software/dendroscope/`
[2] `skelk.sdf-eu.org/terminusest/`

indication that our program is actually able to compute all relevant maximum acyclic agreement forests. Nevertheless, we applied a further test scenario examining this fact in more detail.

4.5.3 Permutation test To check the robustness of our implementation in more detail, we generated a further synthetic dataset containing thousands of tree pairs of low computational complexity, such that each of those tree pairs could be processed by our implementation within less than a minute. More precisely, the dataset contains tree pairs that have been generated in respect to precisely one value for each of the four parameters ℓ, k, c, and p, i.e., $\ell \in \{10\}$, $k \in \{5\}$, $c \in \{1\}$, and $p \in \{30\}$. Next, for each of those tree pairs, we computed two sets of relevant maximum acyclic agreement forests each corresponding to one of both orderings of the two input trees and compared both results.

For each of those tree pairs, both sets of maximum acyclic agreement forests were identical, which means that each maximum acyclic agreement forest that could be computed was always contained in both sets. Notice that by switching the order of the input trees our algorithm runs through different recursive calls, which means that each computational path leading to a maximum acyclic agreement forest usually differs. Nevertheless, due to the fact that the hybridization number is independent from the order of the input trees, those two sets of maximum acyclic agreement forests have to be identical. As we applied this permutation test to thousands of different tree pairs, this is a further strong indication that HYBROSCALE is actually able to compute all relevant maximum acyclic agreement forests for two rooted (nonbinary) phylogenetic $\mathcal{X}$-trees.

4.6 Conclusion

In this section, we have presented the algorithm allMulMAAFs computing a set of relevant maximum acyclic agreement forests for two rooted nonbinary phylogenetic $\mathcal{X}$-trees. allMulMAAFs was developed in respect to the algorithm allHNetworks computing a particular set of minimum hybridization networks for two rooted binary phylogenetic $\mathcal{X}$-trees and is considered to be a first step for making this algorithm accessible to nonbinary phylogenetic $\mathcal{X}$-trees. Additionally, we have established a formal proof showing that the algorithm allMulMAAFs guarantees the computation of all relevant nonbinary maximum acyclic agreement forests.

Moreover, we have integrated our algorithm into the freely available software package HYBROSCALE and, by conducting two specific test scenarios, we have demonstrated the robustness of our implementation. In the next section, we will demonstrate how this algorithm can be used to extend the algorithm allHNetworks so that now minimum hybridization networks displaying the refinements of multiple rooted nonbinary phylogenetic $\mathcal{X}$-trees can be calculated.

5 Discussion

In this work, we have presented the algorithm allMulMAAFs computing a set of relevant maximum acyclic agreement forests for two rooted nonbinary phylogenetic $\mathcal{X}$-trees. allMulMAAFs was developed in respect to the algorithm allHNetworks computing a certain set of minimum hybridization networks for two rooted binary phylogenetic $\mathcal{X}$-trees and is considered to be the first step to make this algorithm accessible to nonbinary phylogenetic $\mathcal{X}$-trees. Additionally, we have provided formal proofs showing that the algorithm allMulMAAFs always guarantees the computation of all relevant nonbinary maximum acyclic agreement forests. Moreover, we have integrated our algorithm into the freely available software package HYBROSCALE and by conducting particular test scenarios, we have demonstrated the robustness of our implementation. It is part of ongoing future work to push on the extension of the algorithm allHNetworks in order to enable the computation of minimum hybridization networks displaying the refinements of multiple rooted nonbinary phylogenetic $\mathcal{X}$-trees.

Bibliography

[1] B. Albrecht, C. Scornavacca, A. Cenci, D. H. Huson, *Fast computation of minimum hybridization networks,* Bioinformatics, 28 (2011), pp. 191–197.

[2] B. Albrecht, *Computing Hybridization Networks for Multiple Rooted Binary Phylogenetic Trees by Maximum Acyclic Agreement Forests,* preprint, arXiv:1408.3044, 2014.

[3] B. Albrecht, *Computing all hybridization networks for multiple binary phylogenetic input trees,* BMC Bioinformatics, **16**:236.

[4] B. Albrecht, *Fast computation of all maximum acyclic agreement forests,* submitted to arXiv.

[5] M. Baroni, S. Gruenewald, V. Moulton, C. Semple, *Bounding the number of hybridisation events for a consisten evolutionary history,* Mathematical Biology, 51 (2005), pp. 171–182.

[6] M. Baroni, C. Semple, *Hybrids in real time,* Syst. Biol., 55 (2006), pp. 46–56.

[7] M. Bordewich, C. Semple, *Computing the Hybridization Number of Two Phylogenetic Trees Is Fixed-Parameter Tractable,* IEEE/ACM Trans. Comput. Biol. Bioinformatics, 4 (2007), pp. 458–466.

[8] M. Bordewich, C. Semple, *Computing the minimum number of hybridization events for a consistent evolutionary history,* Discrete Applied Mathematics, 155 (2007), pp. 914–928.

[9] D. H. Huson, R. Rupp, C. Scornavacca, *Phylogenetic Networks: Concepts, Algorithms and Applications,* Cambridge University Press, 2011.

[10] D. H. Huson, D. Richter, C. Rausch, T. Dezulian, M. Franz, R. Rupp R., *Dendroscope: An interactive viewer for large phylogenetic trees,* BMC Bioinformatics, 8 (2007), pp. 460.

[11] L. J. J. van Iersel, S. Kelk, N. Lekic, N, S. Leen, *Approximation algorithms for nonbinary agreement forests.* SIAM J. Discrete Math., 28 (2014), pp. 49–66.

[12] Z. Li, *Fixed-Parameter Algorithm for Hybridization Number of Two Multifurcating Trees,* Master thesis, Dalhousie University, Halifax, Canada, 2014.

[13] T. Piovesan, S. Kelk, *A simple fixed parameter tractable algorithm for computing the hybridization number of two (not necessarily binary) trees,* preprint, arXiv:1207.6090, 2014.

[14] C. Scornavacca, S. Linz, L., B. Albrecht, *A first step towards computing all hybridization networks for two rooted binary phylogenetic trees,* J. Comp. Biol., 19 (2012), pp. 1227–1243.

[15] C. Whidden, R. Beiko, N. Zeh, *Fixed-Parameter and Approximation Algorithms for Maximum Agreement Forests,* SIAM J. Comput., 42 (2011), 1431–1466.

[16] C. Whidden, R Beiko, N. Zeh, *Fixed-Parameter and Approximation Algorithms for Maximum Agreement Forests of Multifurcating Trees,* preprint, arXiv:1305.0512, 2014.